

UNITED STATES PATENT APPLICATION

FOR

METHOD AND APPARATUS FOR PERFORMING PRIORITY ENCODING
IN A SEGMENTED CLASSIFICATION SYSTEM USING ENABLE SIGNALS

Attorney's Docket No.: NLM.P003

Inventors: Sandeep Khanna

Filed By:

Lawrence M. Cho

P.O. Box 2144

Champaign, IL 61825

(217) 377-2500

EXPRESS MAIL CERTIFICATE OF MAILING

"Express Mail" mailing label number EK 955198639 US

Date of Deposit March 29, 2001

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Assistant Commissioner for Patents, Washington, D. C. 20231

Lawrence M. Cho

(Typed or printed name of person mailing paper or fee)

Lawrence M. Cho

(Signature of person mailing paper or fee)

0981573-03401
TOP SECRET

**METHOD AND APPARATUS FOR PERFORMING PRIORITY ENCODING
IN A SEGMENTED CLASSIFICATION SYSTEM USING ENABLE SIGNALS**
RELATED APPLICATION

This application is a continuation-in-part of, and claims the benefit under
5 Title 35, United States Code, Section 120 of co-pending U.S. Application No.
09/406,170 filed on September 23, 1999 and entitled "METHOD AND
APPARATUS FOR PERFORMING PACKET CLASSIFICATION FOR POLICY-
BASED PACKET ROUTING".

FIELD OF THE INVENTION

10 The present invention relates generally to performing packet classification
for policy-based packet routing.

BACKGROUND

Routers are devices that direct traffic in a network. Each router in the
network has a route table that typically includes routing information to enable
15 incoming packets to be forwarded or routed to their destination in the network.
Some routers forward packets based only on the destination address indicated in
the packet. Other, more complex, routers forward or route packets based on
policies defined, for example, by a network administrator. The latter routing
schemes are commonly referred to as policy-based routing.

20 Policy-based routing can enable packets to be forwarded or routed in a
network based on any number of criteria, including the source of the packet, the
destination of the packet, the cost of forwarding the packet through different
routes or paths in the network, or the bandwidth available on different routes in
the network. Policy-based routing can also be used to provide a certain Quality
25 of Service (QOS) or Type of Service (TOS) to differentiated traffic in the network.
For example, one or more of the various fields (e.g., the TOS bits) in the header

of an Internet Protocol (IP) packet can be used by policy-based routers to forward IP packets in a network.

Each policy-based router implements a policy through the use of route maps that define how to forward the packet in the network. Each route map statement or policy statement contains one or more match clauses and a set clause. The match clauses are a series of conditions that are used to determine if an incoming packet satisfies a particular policy. If all of the match clauses of a policy statement are satisfied, the set clause specifies how the router should forward the packet in the network. If one of the match clauses of a particular policy statement is not satisfied, then the policy-based router investigates subsequent policy statements.

Figure 1 shows exemplary processed policy information 100 of an incoming packet to a policy-based router. Policy information 100 includes several policy fields 102 including a destination address (DA) for the packet, a source address (SA) of the packet, protocol type (PTCL) such as those defined by for an IP packet header, TOS, and COST. Policy information 100 may be received by a policy-based router that implements a policy such as policy 200 shown in Figure 2. Policy 200 includes three separate policy statements 201 through 203. If policy information 100 satisfies the match clause (i.e., the “if” clause) of one of the policy statements, the set clause (i.e., the “then” clause) of that policy statement determines routing information for the packet in the network. For example, if the destination address of the incoming packet is DA1, the source address is SA1, and the TOS field of the packet is TOS1, then routing information RI₂ should be selected.

A policy-based router can use a content addressable memory (CAM)-based system to implement a filtering or classification function to determine whether an incoming packet matches a policy statement. Figure 3 shows one

example of a system 300 that implements policy-based routing using a ternary CAM 302. The policy statements or policy words 201-203 are stored in separate rows in ternary CAM array 304. A ternary CAM array is one that is able to mask entries in a CAM array on a bit-by-bit basis. Ternary CAM array 304 has rows of CAM cells 305 for storing policy field information, and corresponding rows of mask cells 310 for storing mask data. Routing information RI_0 - RI_2 is typically stored in an external memory 308 at addresses corresponding to those at which the respective policy words 201-203 are stored in ternary CAM array 304. Each policy field that corresponds to a match clause for a given policy statement is unmasked by having its corresponding mask bits set, for example, to a logic zero. Conversely, each policy field that does not have a match clause for a given policy statement is masked by having its corresponding mask bits set, for example, to a logic one.

When an incoming packet is received by a policy-based router, it is processed to determine the policy field information. The processed policy field information is provided to system 300 as policy search key 307. For each policy statement in CAM array 304 that matches the policy search key, the corresponding match line ML_0 - ML_N will be asserted and provided to priority encoder 306. In response to the match lines, priority encoder 306 outputs the address of the highest priority matching entry in CAM array 304 to HPM bus 312. If there is more than one matching policy statement in CAM array 304, priority encoder 306 determines that the highest priority matching policy statement is the one stored at the lowest logical address of CAM array 304. For example, as shown in Figure 3, if CAM array 304 is loaded with policy statement 203 at address zero (i.e., the lowest logical address), statement 202 at address one, and statement 201 at address 2, and a policy search key of DA1, SA1, PTCL1, TOS1, COST1 is provided to CAM array 304, then each of policy

statements 201-203 is identified as a match on match lines ML_0 - ML_2 , respectively. Priority encoder 306 outputs address zero on the HPM bus to select route information RI_2 from address zero in memory 308.

Because priority encoder 306 determines the highest priority matching location based on predetermined logical address assignments, policy statements 201-203 are preordered or prioritized such that higher priority policy statements are stored in lower logical addresses of CAM array 304 than lower priority policy statements. A policy statement has a higher priority than another policy statement when the route information for the first policy statement is to be selected over the second policy statement even though both policy statements may match the policy search key (e.g., with masking). The prioritizing of the policy statements is typically performed by table management hardware and/or software, which adds overhead to the router.

If the table is not prioritized correctly, then an incorrect route may be selected for the incoming packet. Figure 4 shows an example in which policy statement 201 is incorrectly prioritized such that it, rather than policy statement 203, is loaded into the lowest logical address. In this example, when policy search key 307 is provided to CAM array 304, priority encoder 306 still provides address zero on HPM bus 312. This causes, however, route information RI_0 to be selected in memory 308 rather than the desired route information RI_2 associated with higher priority policy statement 201. It would be desirable to load policy statements into a CAM array without having to preorder the statements according to their respective priorities.

When a policy is changed by adding a new policy statement that has a higher (or equal) priority than at least one of the policy statements already stored in CAM array 304, the table management hardware and software needs to reprioritize or reorder all or part of CAM array 304. This is typically

accomplished by reloading the CAM array with a new prioritized group of policy statements. This can add significant overhead to the router (e.g., delay and additional hardware and software) to change even just one policy statement in a given policy. If the table is not reloaded, CAM 302 may cause an incorrect route to be selected from memory 308. For example, Figure 5 shows that a new policy statement 204 is loaded into the next free address (i.e., address three) of CAM array 304 without reprioritizing and reloading CAM array 304. Route information RI₃, corresponding to policy statement 204, is loaded into a corresponding address in memory 308. If policy statement 204 has a priority greater than that of 201, 202, or 203, then a search with policy search key 307 will not correctly identify policy statement 204 during a compare operation on CAM array 304. Policy statement 203, rather, will be identified as the higher priority statement, and route information RI₂ will be selected. It would be desirable to be able to add a new policy statement to a policy without reloading the CAM array storing the policy statements.

SUMMARY OF THE INVENTION

A method and apparatus for performing a classification or filtering function for policy-based packet routing in a digital signal processor (DSP) is disclosed. For one embodiment, a digital signal processor includes a first data
5 classification block. The first data classification block outputs a first block priority number associated with a first data stored in the first data classification block that matches a search key. The digital signal processor includes a second data classification block. The second data classification block outputs a second priority number associated with a second data stored in the second data
10 classification block that matches the search key. The digital signal processor includes a device index processor. The device index processor selects a most significant block priority number from the first block priority number and the second block priority number.

Other objects, features, and advantages of the present invention will be
15 apparent from the accompanying drawings and from the detailed description that follows below.

BRIEF DESCRIPTION OF THE DRAWINGS

The features and advantages of the present invention are illustrated by way of example and are by no means intended to limit the scope of the present invention to the particular embodiments shown, and in which:

5 Figure 1 is an example of policy information;
 Figure 2 is an example of a policy having policy statements;
 Figure 3 is an example of storing a policy in a conventional ternary CAM;
 Figure 4 is an example of storing a particular policy in a conventional
ternary CAM and comparing a search key with the policy;

10 Figure 5 is another example of storing a particular policy in a
conventional ternary CAM and comparing a search key with the policy;

 Figure 6 is one embodiment of a classification system for a policy-based
router;

15 Figure 7 is one embodiment of performing the classification operation for
the system of Figure 6;

 Figure 8 is one example of performing the classification operation on a
particular policy stored in the CAM array of Figure 6;

 Figure 9 is one embodiment of reading or writing a policy statement
based on a priority number stored in the priority memory of Figure 6;

20 Figure 10 is one embodiment of the priority index table of Figure 6;

 Figure 11 is another embodiment of the priority index table of Figure 6;

 Figure 12A is one embodiment of two rows of the priority index table of
Figure 11;

25 Figure 12B is another embodiment of two rows of the priority index table
of Figure 11;

 Figure 13 is one example of determining the most significant priority
number stored in the rows of the priority index table;

Figure 14 is one embodiment of a compare circuit and an isolation circuit for the priority logic element of Figure 11;

Figure 15 is another embodiment of a compare circuit and an isolation circuit for the priority logic element of Figure 11;

5 Figure 16 is another example of determining the most significant priority number stored in the rows of the priority index table;

Figure 17 is another embodiment of the priority index table of Figure 6 including inequality circuits;

10 Figure 18 is one example of a policy stored in the digital signal processor of Figure 17;

Figure 19 is one example of loading a new policy statement into the digital signal processor of Figure 18;

Figure 20 is another example of a policy stored in the digital signal processor of Figure 17;

15 Figure 21 is one example of loading a new policy statement into the digital signal processor of Figure 20;

Figure 22 is one embodiment of the rows of the priority index table of Figure 17, in which the priority memory locations are configured as counters;

20 Figure 23 is one embodiment of an inequality circuit and counter of Figure 22;

Figure 24 is one example of determining whether a new number is greater than or equal to priority numbers stored in the priority memory of Figure 17;

Figure 25A is one embodiment of a compare circuit of Figure 24;

Figure 25B is another embodiment of a compare circuit of Figure 24;

25 Figure 26 is one example of determining whether a new number is less than or equal to priority numbers stored in the priority memory of Figure 17;

Figure 27A is one embodiment of a compare circuit of Figure 26;

Figure 27B is another embodiment of a compare circuit of Figure 26;

Figure 28 is one embodiment of a process of deleting or invalidating policy statements and their associated priority numbers from the digital signal processor of Figure 17;

5 Figure 29 is one embodiment of priority index table of Figure 17;

Figure 30 is another embodiment of priority index table of Figure 17;

Figure 31 is one embodiment of an inequality circuit element, a priority logic element and a memory element of Figure 17;

10 Figure 32 is another embodiment of an inequality circuit element, a priority logic element and a memory element of Figure 17;

Figure 33 is another embodiment of the priority index table of Figure 6 including a decrement circuit and storing a policy;

Figure 34 is one example of loading a new policy statement into the digital signal processor of Figure 33;

15 Figure 35 is one embodiment of a row of the priority index table in Figure 33;

Figure 36 is one example, of determining whether a new number is greater than or equal to priority numbers stored in the priority memory of Figure 33;

20 Figure 37 is one embodiment of compare circuit and an isolation circuit for a priority logic element of Figure 35;

Figure 38 is one example, of determining whether a new number is less than or equal to priority numbers stored in the priority memory of Figure 33;

25 Figure 39 is a block diagram of one embodiment of depth cascading digital signal processors;

Figure 40 is a block diagram of one embodiment of a digital signal processor including cascade logic;

Figure 41 is a block diagram of one embodiment of the cascade logic of Figure 39;

Figure 42 is one example of processing Internet protocol addresses based on a classless inter domain routing scheme in the digital signal processor of

5 Figure 6;

Figure 43 illustrates a segmented data classification system according to an embodiment of the present invention;

Figure 44 illustrates a segmented data classification system that includes a device index processor that is an embodiment of the device index processor

10 shown in Figure 43;

Figure 45 illustrates an inhibit signal generator that is an embodiment of the inhibit signal generator shown in Figure 44;

Figure 46 illustrates an address generator that is an embodiment of the address generator shown in Figure 44;

15 Figure 47 is a flow chart illustrating a method for generating a device index in a segmented data classification system according to the present invention;

Figure 48 is a flow chart illustrating a method for generating a device index according to an embodiment of the present invention;

20 Figure 49 illustrates an example of the segmented classification system of Figure 45 generating a device index;

Figure 50 illustrates a segmented data classification system that includes an index processor that is an alternate embodiment of the device index processor shown in Figure 44;

25 Figure 51 illustrates a compare logic that is an embodiment of the compare logic shown in Figure 50;

DETAILED DESCRIPTION

In the following description, for purposes of explanation, specific nomenclature is set forth to provide a thorough understanding of the present invention. However, it will be apparent to one skilled in the art that these
5 specific details may not be required to practice the present invention. In other instances, well known circuits and devices are shown in block diagram form to avoid obscuring the present invention unnecessarily. Additionally, the interconnection between circuit elements or blocks may be shown as buses or as single signal lines. Each of the buses may alternatively be single signal lines,
10 and each of the single signal lines may alternatively be buses. Additionally, the prefix symbol “/” or the suffix “B” attached to signal names indicates that the signal is an active low signal. Each of the active low signals may be changed to active high signals as generally known in the art.

A method and apparatus for performing a classification or filtering
15 function for policy-based packet routing in a digital signal processor (DSP) is disclosed. For one embodiment, the digital signal processor includes a policy statement table for storing policy statements. The policy statement table may be stored in a ternary CAM array that stores the policy statements and associated mask data. Each policy statement has associated with it a priority number that
20 indicates the priority of the policy statement relative to other policy statements. The priority numbers are separately stored in a priority index table. The priority index table includes priority logic that determines the most significant priority number from among the policy statements that match an incoming packet during a classification or filter operation. The priority logic also identifies the
25 location in the priority index table of the most significant priority number. The most significant priority number may be the priority number with the lowest or highest numerical value. The identified location in the priority index table can be

used to access associated route information or other information stored in a route memory array located external to, or on the DSP. When the route memory array is external to the DSP, the DSP may include an encoder to encode the identified location in the priority index table into an address for the route memory.

5 The DSP configuration obviates preloading the policy statements in the policy statement table in a predetermined order. Instead, the priority logic determines the most significant priority number from among matching policy statements regardless of the order in which the policy statements are stored in the table. This can reduce the hardware and/or software needed for table
10 management of the table, and can increase the performance of a router incorporating the DSP.

 In addition, new policy statements can be added at any location in the policy statement table, and associated priority numbers loaded into corresponding locations in the priority index table. If a new policy statement has
15 a priority that is greater than or equal to a priority of a policy statement already stored in the policy statement table, the priority number of the previously stored policy statement may be updated to accommodate the new policy statement. Similarly, when a policy statement is removed (i.e., invalidated or overwritten) from the policy statement table, the priority numbers of the previously stored
20 lower priority policy statements may be updated. The updating functions can be performed by the priority logic in the priority index table, or by inequality circuits in the priority index table. The updating functions can be performed without the need to physically reorder the policy statements in the policy statement table, or to physically reorder the priority numbers in the priority
25 index table. This also can reduce the hardware and/or software needed for table management of the policy statement table, and can increase the performance of a router incorporating the DSP.

The DSP can also be used in other non-networking applications. For example, the DSP can be used to process if-then-else functions in other applications.

Classifying or Filtering Policy Statements

5 Figure 6 shows one embodiment of a classification or filtering system 600 for a policy-based router. System 600 includes DSP 602 and route memory 614. DSP 602 includes policy statement table 604, priority index table 606 and encoder 612, and may be an integrated circuit component formed on a single semiconductor substrate.

10 For this embodiment, policy statement table 604 is stored in a ternary CAM array that stores policy statements or policy words. Ternary CAM array 604 has rows of CAM cells 605 for storing policy field information PF1-PFX, where X is any number. Each policy field PF1-PFX can include any policy information including DA, SA, PTCL, TOS, and COST, or any other type of
15 policy field to assist in the classification or filtering of the policy statement to provide a certain Quality of Service (QoS), Class of Service (CoS), and the like. Each policy field may include any number of bits. Additional information associated with each policy field may be stored in one or more additional binary or ternary CAM cells or other types of memory cells disposed in each row of
20 CAM 604. Ternary CAM 604 also has rows of mask cells 607 for storing mask data M1-MX corresponding to each row of policy fields 605. Global masks (not shown) may be used to mask entire columns in CAM array 604 as generally known in the art. For alternative embodiments, CAM array 604 may be any other type of CAM including a binary CAM, or any other type of memory to
25 store policy statements to be compared with processed policy information of an incoming packet.

0981578-032401

Ternary CAM array 604 can be any ternary CAM array that is capable of locally masking each entry on a bit-by-bit basis. Each policy field that corresponds to a match clause for a given policy statement will be unmasked by having its corresponding mask bits set, for example, to a logic zero (or, alternatively, a logic one). Conversely, each policy field that does not have a match clause for a given policy statement will be masked by having its corresponding mask bits set, for example, to a logic one (or, alternatively, a logic zero). As each policy field, and/or individual bits within a policy field, can be masked, DSP 602 also supports rule and route aggregation. That is, DSP 602 supports ranges of addresses or policy field information.

Each policy statement loaded into ternary CAM array 604 has associated with it a priority number P_0 - P_Z and route information RI_0 - RI_{N-1} . The priority number indicates the priority of the policy statement relative to other policy statements in a given policy. The policy numbers may be assigned by a user of DSP 602 including, for example, a network administrator or the router itself. The priority numbers P_0 - P_Z are separately stored at locations 608_0 - 608_{N-1} , respectively, of priority memory 608 of priority index table 606. Route information RI_0 - RI_{N-1} for the particular policy statements are stored at locations 614_0 - 614_{N-1} , respectively, in route memory 614. The route information may include, for example, forwarding or next hop information, authentication information, QOS, TOS, time to live information or other packet filtering and classification information for an incoming packet to the router incorporating system 600. A policy statement, its priority number, and its route information are each stored at the corresponding addresses in each of their respective memory arrays.

Priority memory 608 and route memory 614 may each be any type of memory array including volatile, non-volatile, random access memory (RAM),

and/or read only access memory (ROM). For one embodiment, priority memory 608 comprises a CAM array.

Priority memory 608 may be n bits wide to accommodate $Z=2^n$ priority numbers, where n is any number. The number of addressable locations N in priority memory 608 may be greater than, less than, or equal to Z . For one example, priority memory 608 may be 20 bits wide to accommodate up to 2^{20} or 1 Meg (i.e., 1,048,576) priority numbers.

The priority numbers may be assigned in ascending priority order such that zero is the highest priority number and 2^n-1 is the lowest priority number.

Alternatively, the priority numbers may be assigned in descending priority order such that 2^n-1 is the highest priority number and zero is the lowest priority number. Each priority number may be assigned so as to identify the priority of each policy statement relative to other policy statements. For one embodiment, the priority numbers may be assigned consecutively. For example, the highest priority policy statement can be assigned the highest priority number (e.g., zero or 2^n-1), the next lower priority policy statement can be assigned the next lower priority number (e.g., one or 2^n-2), and so forth. For another embodiment, gaps may be left in the priority number assignments to allow for the addition of future priority numbers associated with new policy statements.

Priority index table 606 also includes priority logic 610 that compares the priority numbers with each other for all corresponding policy statements that match an incoming packet. Priority logic 610 identifies the most significant priority number PNUM in memory 608 from among the compared priority numbers, and further identifies the location of PNUM in priority memory 608.

PNUM has the lowest numerical value when the priority numbers are assigned in ascending priority order, and PNUM has the highest numerical value when the priority numbers are assigned in descending priority order. Priority logic

610 may also output PNUM from DSP 602. The identified location of PNUM in the priority memory is provided on internal address lines IAD₀-IAD_{N-1} to encoder 612. For one example, one of IAD₀-IAD_{N-1} is asserted to indicate the location in priority memory 608 of PNUM from among the compared priority numbers. This location also corresponds to the location of the highest priority matching policy statement in ternary CAM array 604.

The address of the identified location of the highest priority matching policy statement in ternary CAM array 604 is determined by encoder 612 and output to HPM bus 616. The encoded address can then be used to access the corresponding route information from memory 614. Encoder 612 may be any encoding logic that takes the information on address lines IAD₀-IAD_{N-1} and generates an encoded address. For one embodiment, encoder 612 is a ROM.

For another embodiment, route memory 614 may also be included within DSP 602. For this embodiment, encoder 612 may be omitted and route memory 614 may be accessed directly by internal address lines IAD₀-IAD_{N-1}.

Figure 7 summarizes the classification or filtering function 700 (i.e., search or compare operation) performed by DSP 602 for an incoming packet according to a policy stored in ternary CAM 604. An incoming packet received by a policy-based router incorporating system 600 is initially processed to determine the policy field information. The policy field information is provided to system 600 as policy search key 609. At step 702, the policy fields of policy search key 609 are compared with the policy statements stored in ternary CAM array 604. For each policy statement that matches the policy search key, the corresponding match line ML₀-ML_{N-1} is asserted. If no match is found, then the process stops at step 704.

At step 706, priority logic 610 determines PNUM and identifies its location in priority memory 608. The identified location is provided on internal

address lines IAD_0 - IAD_{N-1} to encoder 612. At step 708, encoder 612 determines the address of the identified location in priority index table 606. This encoded address is also the logical address of the highest priority matching policy statement in ternary CAM array 604. Encoder 612 outputs the encoded address to HPM bus 616. The encoded address can then be used at step 710 to access the corresponding route information in memory 614. Steps 708 and/or 710 may be omitted when encoder 612 is removed from DSP 602, and priority logic 610 may directly access the route information in memory 614.

For another embodiment, IAD_0 - IAD_{N-1} are provided to CAM array 604 to access the highest priority matching policy statement, which may then be read from DSP 602. Alternatively, HPM bus 616 may be provided to CAM array 604 (e.g., through a decoder) to access the highest priority matching policy statement.

Figure 8 shows one example in which policy statements 201, 202, and 203 from Figure 2 are stored at locations 604_0 , 604_1 , and 604_2 (i.e., addresses 0, 1, and 2), respectively, of ternary CAM array 604. The corresponding priority numbers 2, 1, and 0 are stored at locations 608_0 , 608_1 , and 608_2 , respectively, in memory 608. Additionally, the corresponding route information RI_0 , RI_1 , and RI_2 are stored at locations 614_0 , 614_1 , and 614_2 , respectively of route memory 614. The policy statements and priority numbers are written into their respective memories using conventional write circuits, counters, and/or address decoders, etc. (not shown).

For this embodiment, the priority numbers have been assigned in ascending priority order such that policy statement 203 is identified as the highest priority policy statement by being assigned priority number 0, the lowest numerical value; policy statement 201 is identified as the having the lowest priority policy statement by being assigned priority number 2, the highest

numerical value; and, policy statement 202 is identified as having a priority greater than that of policy statement 201, but less than that of policy statement 203, by being assigned priority number 1.

For an alternative embodiment, the priority numbers may be assigned in descending priority order such that policy statement 201 is identified as the highest priority policy statement by being assigned priority number 2, the highest numerical value; policy statement 203 is identified as having the lowest priority policy statement by being assigned priority number 0, the lowest numerical value; and, policy statement 202 is identified as having a priority greater than that of policy statement 201, but less than that of policy statement 203, by being assigned priority number 1.

The process of determining the route information for policy search key 609 is illustrated with the aid of Figure 7. At step 702, the policy fields of policy search key 609 are compared with the policy statements stored in ternary CAM array 604. In this example, the policy search key has policy fields of DA=DA1, SA=SA1, PTCL=PTCL1, TOS=TOS1, and COST=COST1. CAM array 604 determines that each of the policy statements 201-203, as masked by their respective mask data, matches policy search key 609. In response, each of match lines ML₀-ML₂ is asserted.

At step 706, priority logic 610 compares, with each other, priority numbers 0, 1, and 2 associated with matching policy statements 203, 202, and 201 respectively. Priority logic 610 determines that priority number 0 is the most significant priority number, asserts IAD₂, and de-asserts the other internal address lines. Encoder 612 encodes the internal address information, at step 708, and generates an external address of 2 on HPM bus 616. The external address can be used to access route information RI₂ stored at address two in route memory 614.

In contrast to the conventional system shown in Figure 4, DSP 602 is able to identify the highest priority matching policy statement stored in ternary CAM array 604 regardless of where the policy statements are stored in CAM array 604.

The process illustrated in Figure 7 identifies the location in priority memory 608 of the most significant priority number from among the compared priority numbers. Once this location is identified, the priority number stored at the identified location can be read out from DSP 602 by a read circuit (not shown), or a new priority number can be written into that location by a write circuit (not shown). In the former case, the user of the DSP (e.g., a network administrator or the policy-based router itself) can determine what priorities have been assigned to policy statements already stored in CAM array 604. In the latter case, priority numbers can be updated by the user for already stored policy statements. This provides the user with flexibility in the control and management of the policy statements stored in DSP 602.

For added flexibility, the user can read a policy statement (e.g., one or more of the policy fields and/or one or more of the corresponding mask data) based on priority numbers already stored in the priority memory, or write a new policy statement for a priority number already stored in the priority memory. For these embodiments, priority memory 608 may be a CAM. For an alternative embodiment, each priority number may be separately compared by a comparison circuit with an externally applied priority number.

Figure 9 shows one embodiment of a process 900 for reading or writing a policy statement based on a priority number already stored in the priority memory. At step 902, an external priority number is compared with the priority numbers stored in priority memory 608. The valid priority numbers stored in memory 608 may be indicated by one or more validity bits (not shown) stored in CAM memory 604 or memory 608. If there is no match, or more than one match

(i.e., a multiple match), then the process stops at step 904. Alternatively, a priority encoder circuit can be coupled to priority index table 606 to select one of the multiple matches for step 906.

If there is a single match, then an indication of the matching location is provided on lines IAD_0 - IAD_{N-1} to encoder 612 at step 906. When priority memory 608 is implemented as a CAM, IAD_0 - IAD_{N-1} may be the match lines associated with each priority number storage location. At step 908, encoder 612 determines the address of the matching location in priority memory 608 and provides this address on HPM bus 616. This address is also the address of the corresponding policy statement in CAM array 604. At step 910, the encoded address can then be provided to CAM array 604 (e.g., through a decoder), and the policy statement stored at that address either read from DSP 602, or overwritten. For an alternative embodiment, IAD_0 - IAD_{N-1} may be provided directly to CAM array 604 to access the desired policy statement and the process may omit step 908. Additionally, a new priority number can be written to priority memory 608 at the address determined in step 910.

With reference again to Figures 6 and 7, priority index table 606 stores priority numbers in priority memory 608, and priority logic 610 determines the most significant priority number among those associated with policy statements that match policy search key 609. Figure 10 shows DSP 1000 that includes priority index table 1001 that is one embodiment of priority index table 606. In this embodiment, CAM array 1002 is one embodiment of priority memory 608, and priority logic 1004 is one embodiment of priority logic 610.

DSP 1000 performs the filtering function of Figure 7 as follows. At step 702, a policy search key is compared with the policy fields stored in CAM array 604 and match lines ML_0 - ML_{N-1} are asserted for the matching locations. The asserted match lines access the associated priority numbers stored in CAM array

1002. For one embodiment, ML_0 - ML_{N-1} are coupled to corresponding word lines of the rows of CAM cells in CAM array 1002. In response to the asserted match lines, CAM array 1002 provides the selected priority numbers to priority logic 1004 over bus 1006 to determine which priority number is the most significant priority number. The priority numbers may be provided one at a time over bus 1006, or in parallel.

Priority logic 1004 includes one or more compare circuits to compare the priority numbers and determine the most significant priority number. For one embodiment, priority logic 1004 includes a compare circuit and a register. The first priority number provided on bus 1006 is stored in the register and compared with the second priority number. Each subsequent priority number provided over bus 1006 is then compared and the result (either greater than or less than) is then stored in the register. This process continues until no more priority numbers are provided over bus 1006. For another embodiment, multiple compare circuits may be provided to compare multiple priority numbers at the same time to generate the most significant priority number.

Once the most significant priority number is determined, it is provided over bus 1008 to CAM 1002 and searched against the stored priority numbers to determine where it is located in CAM array 1002. The indicated location is provided on internal address lines IAD_0 - IAD_{N-1} , which may correspond to the match lines of CAM array 1002. Steps 708 and 710 may then be performed as previously described above.

Figure 11 shows DSP 1100 that includes priority index table 1101 that is another embodiment of priority index table 606. In this embodiment, priority memory 608 and priority logic 610 are merged together on a bit-by-bit basis to form priority index table 1101. The priority memory includes memory elements 1102 that each store one bit of a priority number for a given row. Each memory

element may be any type of storage mechanism including volatile or non-volatile memory cells. The priority logic includes priority logic elements 1104. Each priority logic element 1104 is associated with, or corresponds to, one of the memory elements 1102 such that columns 1106₀-1106_{n-1} of priority index table 1101 have a memory element/ priority logic element pair for each of its rows. Each priority logic element 1104 effectively compares the priority number bit stored in its associated memory element 1102 with the priority number bits stored in every other memory element of its column to determine one of bits PNUM₀-PNUM_{n-1} for the most significant priority number. Bits PNUM₀-PNUM_{n-1} comprise the most significant priority number from among the policy statements that match a policy search key.

Figure 12A shows priority index table 1200 that is one embodiment of two rows of priority index table 1101. For this embodiment, each priority logic element 1104 includes a compare circuit 1206 and an isolation circuit 1204. Each compare circuit 1206 is connected in a wired-OR configuration with the other compare circuits in its respective column by one of priority signal lines 1208₀-1208_{n-1}. Each priority signal line may be pre-charged towards a power supply voltage (or any other predetermined voltage) by a pre-charge circuit 1202. Each compare circuit 1206 may be any digital or analog compare circuit that, when executing step 706 of Figure 7, effectively compares the priority number bit stored in its respective storage element 1102 with the priority number bits stored in every other storage element 1102 of the same column. Additionally, each compare circuit monitors the comparison result of the more significant priority number bits through the logical states of match line segments 1210. Match line segments 1210 are coupled between match lines ML₀-ML_{N-1} and internal address lines IAD₀-IAD_{N-1} by isolation circuits 1204. The isolation circuits isolate the comparison results generated for less significant priority bit locations from

affecting the comparison results generated for more significant priority bit locations. The isolation circuits may also work together with the comparison circuits to control the state of the match line segments.

The operation of priority index table 1200 can be illustrated with an example shown in Figure 13 and with the aid of Figure 7. In this example, priority index table 1200 comprises a 2x4 matrix of rows and columns. For other embodiments, any numbers of rows and columns can be used. Row zero stores priority number 0110 having the decimal equivalent of the number 6, and row one stores priority number 0101 having the decimal equivalent of the number 5.

For this example, each of row zero and row one of CAM array 604 have policy statements that match the policy search key such that match lines ML_0 and ML_1 are asserted (step 702). Also, for this example, the priority numbers are stored in ascending priority order such that 0101 is the more significant priority number between 0101 and 0110.

At step 706, compare circuits 1206_{0,0}-1206_{3,1} determine that 0101 is the more significant priority number PNUM, and cause IAD_1 to be asserted to indicate that 0101 is stored in row one of the priority index table. Compare circuits 1206_{0,0}-1206_{3,1} determine that PNUM is 0101 as follows. The most significant bit PNUM₃ is resolved first. When any memory element 1102 stores a logic zero and the corresponding match line segment 1210 is asserted, the corresponding priority signal line 1208 is discharged. Thus, each of compare circuits 1206_{3,1} and 1206_{3,0} discharge signal line 1208₃ such that PNUM₃ is a logic zero. Additionally, compare circuit 1206_{3,1} compares the state of priority signal line 1208₃ with the priority number bit stored in 1102_{3,1}, and determines that both have the same logic state. This causes compare circuit 1206_{3,1} not to affect the logical state of match line segment 1210_{2,1} such that match line segment 1210_{2,1} has the same logic state as match line segment 1210_{3,1} (ML_1). Similarly, compare

circuit 1206_{3,0} compares the state of priority signal line 1208₃ with the priority number bit stored in 1102_{3,0} and determines that both have the same state. This causes compare circuit 1206_{3,0} not to affect the logical state of match line segment 1210_{2,0} such that match line segment 1210_{2,0} has the same logic state as match line segment 1210_{3,0} (ML₀).

The next most significant bit PNUM₂ is then resolved. Memory elements 1102 that store a logic one do not discharge their corresponding priority signal lines 1208. Since memory elements 1102_{2,1} and 1102_{2,0} both store logic one states, signal line 1208₂ remains pre-charged such that PNUM₂ is a logic one.

Additionally, compare circuit 1206_{2,1} compares the state of priority signal line 1208₂ with the priority number bit stored in 1102_{2,1}, and determines that both have the same logic state. This causes compare circuit 1206_{2,1} not to affect the logical state of match line segment 1210_{1,1} such that match line segment 1210_{1,1} has the same logic state as match line segment 1210_{2,1}. Similarly, compare circuit 1206_{2,0} compares the state of priority signal line 1208₂ with the priority number bit stored in 1102_{2,0} and determines that both have the same logic state. This causes compare circuit 1206_{2,0} to not affect the logical state of match line segment 1210_{1,0} such that match line segment 1210_{1,0} has the same logic state as match line segment 1210_{2,0}.

PNUM₁ is resolved next. Since memory element 1102_{1,1} stores a logic zero and match line segment 1210_{1,1} is asserted, compare circuit 1206_{1,1} discharges priority signal line 1208₁. This causes PNUM₁ to be a logic zero. Additionally, compare circuit 1206_{1,1} compares the logic zero state of priority signal line 1208₁ with the logic zero stored in 1102_{1,1} and allows match line segment 1210_{0,1} to have the same state as match line segment 1210_{1,1}. Compare circuit 1206_{1,0}, however, compares the logic zero on priority signal line 1208₁ with the logic one stored in memory element 1102_{1,0}, and de-asserts match line segment 1210_{0,0}.

When a match line segment is de-asserted, all subsequent compare circuits for that row will de-assert the remaining match line segments of the row such that the corresponding internal address line IAD will be de-asserted. When IAD is de-asserted for a particular row, this indicates that the most significant priority number is not stored in that row. Additionally, when the remaining match line segments are de-asserted for a row, the compare circuits for that row do not discharge the remaining priority signal lines regardless of the logic states stored in the corresponding memory elements of that row. For example, compare circuit 1206_{0,0} does not discharge priority signal line 1208₀ even though memory element 1102_{0,0} stores a logic zero. Additionally, isolation circuits 1204_{3,0}, 1204_{2,0}, and 1204_{1,0} isolate the de-asserted match line segment 1210_{0,0} from match line segment 1210_{3,0}, 1210_{2,0}, and 1210_{1,0} such that PNUM₃, PNUM₂, and PNUM₁ are not affected by the de-assertion of match line segment 1210_{0,0}.

Lastly, the least significant bit PNUM₀ is resolved. Compare circuit 1206_{0,1} alone determines PNUM₀ since compare circuit 1206_{0,0} can not discharge priority signal line 1208₀. Since memory element 1102_{0,1} stores a logic one and match line segment 1210_{0,1} is asserted, compare circuit 1206_{0,1} leaves priority signal line 1208₀ pre-charged, and PNUM₀ is a logic one. Additionally, compare circuit 1206_{0,1} allows IAD₁ to have the same state as match line segment 1210_{0,1}. Since match line segment 1210_{0,1} is asserted, IAD₁ will be asserted indicating that the most significant priority number is stored in that row.

Thus, when the processing of step 706 is completed, bits PNUM₃-PNUM₀ indicate that the most significant priority number stored in the priority index table is 0101, and IAD₁ is asserted identifying that 0101 is stored in row one.

Any circuits may be used for compare circuits 1206 and/or isolation circuits 1204 to implement the process illustrated above. Table 1 shows one example of a truth table for implementing each compare circuit 1206, where X

(column) and Y (row) are any integers. Other truth tables may be used (and corresponding logic generated accordingly) including those that logically complement one or more of the signals indicated in Table 1.

| STATE | 1208 | 1102 | 1210 _{XY} | 1210 _{X-1,Y} |
|-------|------|------|--------------------|-----------------------|
| 0 | 0 | 0 | 0 | 0 |
| 1 | 0 | 0 | 1 | 1 |
| 2 | 0 | 1 | 0 | 0 |
| 3 | 0 | 1 | 1 | 0 |
| 4 | 1 | 0 | 0 | 0 |
| 5 | 1 | 0 | 1 | 1 |
| 6 | 1 | 1 | 0 | 0 |
| 7 | 1 | 1 | 1 | 1 |

Table 1

Note that state 5 may not occur since priority signal line 1208 will not be a logic one when a memory element 1102 stores a logic zero and the corresponding match line segment 1210 is asserted to a logic one state. For other truth tables, state 5 may occur.

Any logic or circuitry may be used to implement the truth table of Table 1. Figure 14 shows one embodiment of compare circuit 1206_{n-1,0} and isolation circuit 1204_{n-1,0} for implementing the truth table of Table 1. Compare circuit 1403 is one embodiment of compare circuit 1206_{n-1,0}, and isolation circuit 1401 is one embodiment of isolation circuit 1204_{n-1,0}. The embodiment of Figure 14 may be used to implement all of the priority logic elements 1104 in the priority index table.

Compare circuit 1403 includes inverter 1414, transistors 1406 and 1408 connected in series between priority signal line 1208_{n-1} and ground, and

transistors 1410 and 1412 connected in series between match line segment 1210_{n-2,0} and ground. N-channel transistor 1406 has its drain coupled to signal line 1208_{n-1}, its gate coupled to match line segment 1210_{n-1,0}, and its source coupled to the drain of n-channel transistor 1408. Transistor 1408 has its gate coupled to receive the logical complement of the priority number bit (/D) stored in memory element 1102_{n-1,0}, and its source coupled to ground. N-channel transistor 1410 has its drain coupled to match line segment 1210_{n-2,0}, its gate coupled to signal line 1208_{n-1} via inverter 1414, and its source coupled to the drain of n-channel transistor 1412. Transistor 1412 has its gate coupled to receive the priority number bit (D) stored in memory element 1102_{n-1,0}, and its source coupled to ground. Any of transistors 1406, 1408, 1410, and 1412 can be replaced with other types of transistors and the logic adjusted accordingly.

Isolation circuit 1401 includes inverters 1402 and 1404. For alternative embodiments, only one inverter may be used and the logic of the next compare circuit adjusted accordingly. For other embodiments, other isolation circuits such as one or more AND, OR, or XOR logic gates or pass gates may be used.

Figure 15 shows another embodiment of compare circuit 1206_{n-1,0} and isolation circuit 1204_{n-1,0} for implementing the truth table of Table 1. In this embodiment, the compare circuit and isolation circuit are merged into the same logic that includes transistors 1406 and 1408 configured as in Figure 14, inverter 1501, NAND gate 1502, and AND gate 1503. NAND gate 1502 has one input coupled to signal line 1208_{n-1} via inverter 1501, and another input coupled to receive the priority number bit (D) stored in memory element 1102_{n-1,0}. AND gate 1503 has one input coupled to match line segment 1210_{n-1,0}, another input coupled to the output of NAND gate 1502, and an output coupled to match line segment 1210_{n-2,0}.

For the example described above with respect to Figure 13, the most significant priority number is the lowest number such that 0101 is the most significant number between 0101 and 0110. For another embodiment, the priority numbers are stored in descending priority order such that 0110 is the most significant priority number between 0101 and 0110. For this embodiment, compare circuits 1206_{0,0}-1206_{3,1} determine that 0110 is the most significant priority number, and assert IAD₀ to indicate that 0110 is stored in row zero of the priority index table. For this embodiment, as shown in Figure 16, inverters 1212₃-1212₀ output the logical states of priority signal lines 1208₃-1208₀, respectively, as PNUM₃-PNUM₀. For other embodiments, inverters 1212₃-1212₀ are not required. Compare circuits 1206_{0,0}-1206_{3,1} and isolation circuits 1204_{0,0}-1204_{3,1} determine that 0110 is the highest priority number as follows.

As in the previous embodiment, the most significant bit PNUM₃ is resolved first. For this embodiment, memory elements 1102 that store a logic zero do not discharge their corresponding priority signal lines 1208. Since both of memory elements 1102_{3,1} and 1102_{3,0} store logic zero states, signal line 1208₃ remains pre-charged such that PNUM₃ is a logic zero. Additionally, compare circuit 1206_{3,1} compares the state of priority signal line 1208₃ with the priority number bit stored in 1102_{3,1} and determines that both have different logical states. This causes compare circuit 1102_{3,1} to not affect the logical state of match line segment 1210_{2,1} such that match line segment 1210_{2,1} has the same logic state as match line segment 1210_{3,1}. Similarly, compare circuit 1206_{3,0} compares the state of priority signal line 1208₃ with the priority number bit stored in 1102_{3,0} and determines that both have the different logical states. This causes compare circuit 1102_{3,0} to not affect the logical state of match line segment 1210_{2,0} such that match line segment 1210_{2,0} has the same logic state as match line segment 1210_{3,0}.

PNUM₂ is resolved next. When any memory element 1102 stores a logic one and the corresponding match line segment 1210 is asserted, the corresponding priority signal line 1208 is discharged. Thus, each of memory compare circuits 1206_{2,1} and 1206_{2,0} discharge signal line 1208₂ such that PNUM₂ is a logic one. Additionally, compare circuit 1206_{2,1} does not affect match line segment 1210_{1,1} because memory element 1102_{2,1} and priority signal line 1208₂ have different logic states. Thus, match line segment 1210_{1,1} will have the same logic state as match line segment 1210_{2,1}. Similarly, compare circuit 1206_{2,0} does not affect match line segment 1210_{1,0} because memory element 1102_{2,0} and priority signal line 1208₂ have different logic states. Thus, match line segment 1210_{1,0} has the same state as match line segment 1210_{2,0}.

PNUM₁ is resolved next. Since memory element 1102_{1,0} stores a logic one and match line segment 1210_{1,0} is asserted, compare circuit 1206_{1,0} discharges priority signal line 1208₁. This causes PNUM₁ to be a logic one. Additionally, compare circuit 1206_{1,0} compares the logic zero state of priority signal line 1208₁ with the logic one stored in 1102_{1,0} and allows match line segment 1210_{0,0} to have the same logic state as match line segment 1210_{1,0}. Compare circuit 1206_{1,1}, however, compares the logic zero on priority signal line 1208₁ with the logic zero stored in memory element 1102_{1,1}, and de-asserts match line segment 1210_{0,1}. As in the example of Figure 13, when a match line segment is de-asserted, all subsequent compare circuits for that row will de-assert the remaining match line segments for the row such that the corresponding internal address line IAD will be de-asserted. When IAD is de-asserted for a particular row, this indicates that the most significant priority number is not stored in that row. Additionally, when the remaining match line segments are de-asserted, the compare circuits for that row do not discharge the remaining priority signal lines regardless of the logic states stored in their corresponding memory elements of that row. For

example, compare circuit 1206_{0,1} does not discharge priority signal line 1208₀ even though memory element 1102_{0,1} stores a logic one. Additionally, isolation circuits 1204_{3,1}, 1204_{2,1}, and 1204_{1,1} isolate the de-asserted match line segment 1210_{0,1} from match line segments 1210_{3,1}, 1210_{2,1}, and 1210_{1,1} such that the PNUM₃, PNUM₂, PNUM₁ and are not affected by de-asserted match line segment 1210_{0,1}.

Lastly, the least significant priority number bit PNUM₀ is resolved.

Compare circuit 1206_{0,0} alone determines PNUM₀ since compare circuit 1206_{0,1} can not discharge priority signal line 1208₀. Since memory element 1102_{0,0} stores a logic zero and match line segment 1210_{0,0} is asserted, compare circuit 1206_{0,0}

leaves priority signal line 1208₀ pre-charged, and PNUM₀ is a logic zero.

Additionally, compare circuit 1206_{0,0} allows IAD₀ to have the same logic state as match line segment 1210_{0,0}. Since match line segment 1210_{0,0} is asserted, IAD₀ will be asserted indicating that the most significant priority number is stored in that row.

Thus, when the processing of step 706 is completed, PNUM₃-PNUM₀, for this embodiment, indicate that the most significant priority number stored in the priority index table is 0110, and IAD₀ is asserted identifying that 0110 is stored in row zero.

Any circuits may be used for compare circuits 1206 and/or isolation circuits 1204 to implement the process illustrated above for Figure 16. Table 2 shows one example of a truth table for implementing each compare circuit 1206, where X (column) and Y (row) are any integers. Other truth tables may be used (and corresponding logic generated accordingly) including those that logically complement one of more of the signals indicated in Table 2.

| STATE | PNUM | 1102 | 1210 _{XY} | 1210 _{X-1,Y} |
|-------|------|------|--------------------|-----------------------|
| 0 | 0 | 0 | 0 | 0 |
| 1 | 0 | 0 | 1 | 1 |
| 2 | 0 | 1 | 0 | 0 |
| 3 | 0 | 1 | 1 | 1 |
| 4 | 1 | 0 | 0 | 0 |
| 5 | 1 | 0 | 1 | 0 |
| 6 | 1 | 1 | 0 | 0 |
| 7 | 1 | 1 | 1 | 1 |

Table 2

Note that state 3 may not occur since PNUM will not be a logic zero when a memory element 1102 stores a logic one and the corresponding match line segment 1210 is asserted to a logic one state. For other truth tables, state 3 may occur.

Any logic or circuitry may be used to implement the truth table of Table 2. For one embodiment, the compare circuit 1403 and isolation circuit 1401 of Figure 14 may be used. For this embodiment, transistor 1408 has its gate coupled to receive the priority number bit (D) stored in memory element 1102_{n-1,0}, and transistor 1412 has its gate coupled to receive the logical complement of the priority number bit (/D) stored in memory element 1102_{n-1,0}. Additionally, signal line 1208_{n-1} provides the logical complement of PNUM_{n-1}.

For another embodiment, the logic and circuitry of Figure 15 may be used to implement the truth table of Table 2. For this embodiment, transistor 1408 has its gate coupled to receive the priority number bit (D) stored in memory element 1102_{n-1,0}, and NAND gate 1502 has an input coupled to receive the logical complement of the priority number bit (/D) stored in memory element 1102_{n-1,0}. Additionally, signal line 1208_{n-1} provides the logical complement of PNUM_{n-1}.

Figure 12B shows another embodiment of 1201₀ of the priority index table. For this embodiment, IAD₀ is determined on signal line 1216 that is separate from the match line segments. As shown in Figure 12B, each match line signal line is coupled to a discharge circuit that includes an inverter-transistor pair (i.e., inverters 1212_{n-1,0}-1212_{0,0} and corresponding transistors 1214_{n-1,0}-1214_{0,0}) such that if any of the match line segments is discharged, then signal line 1216 is discharged to a low logic state. For other embodiments, other discharge circuitry may be used. Signal line 1216 is pre-charged to a high logic state by pre-charge circuit 1202. For an alternative embodiment, signal line 1216 may be driven directly by the match line segments.

The previously described embodiments of system 600 of Figure 6 show that policy statements can be loaded into CAM array 604 in any order. When an incoming packet is received, DSP 602 can identify the address in CAM array 604 of the highest priority policy statement that matches the policy information of the incoming packet. The identified address can then be used to access routing information stored in route memory 614. DSP 602 can perform this function without the user having to preorder the policy statements for entry into the CAM array. This can reduce the hardware and/or software needed for table management of the CAM array, and can increase the performance of a router incorporating the DSP.

DSP 602 can operate asynchronously or synchronously. When DSP 602 operates synchronously, it receives a clock signal that may be used to clock in the policy search key and an instruction that causes the process of Figure 7 to be performed by DSP 602. DSP 602 may implement the classification function of Figure 7 in one or more clock cycles.

Inserting and Deleting Policy Statements

As previously described, priority numbers for policy statements may be assigned in ascending or descending priority order such that there are gaps left between the numbers to accommodate the new priority numbers associated with new policy statements to be stored in CAM array 604. Alternatively, the priority numbers may be assigned in consecutive ascending or descending priority order. New policy statements and their associated priority numbers can be added to the tables 604 and 608 in conformance with either assignment method without having to reload or physically reorder the policy statements or the priority numbers in the respective tables.

Each new policy statement can be loaded into any location (e.g., the next free location) in CAM array 604, and can be assigned a new priority number without having to reload or reorder CAM array 604 and priority memory 608. When a new policy statement is received, its priority number can be compared with the existing priority numbers already stored in priority memory 608 to determine if a policy statement already exists that has been assigned that priority. It is generally desirable that no two policy statements have the same priority number. Thus, if the priority number already exists, the network administrator or the policy-based router itself can assign the new policy statement a new priority number, or the priority number of the existing policy statement can be updated (i.e., incremented or decremented). Since the existing priority numbers are stored in ascending or descending order, updating one priority number may also result in the need to update other priority numbers such that no two priority numbers are the same.

For one embodiment, DSP 602 may include additional circuitry to determine if at least one of the existing priority numbers stored in memory 608 is greater than or equal to (e.g., for ascending priority order), or, alternatively, less than or equal to (e.g., for descending priority order), the new priority number. If

so, the existing priority numbers that are identified by the comparison may be updated such that a new policy statement does not have the same priority number as an existing policy statement. For other embodiments described below, priority logic 610 itself in priority index table 606 may perform the additional compare function(s) and control the updating function(s).

Figure 17 shows one embodiment of priority index table 606 of Figure 6 including inequality circuits INEQ 1704₀-1704_{N-1} that perform the additional comparison functions and control the updating functions. The inequality circuits determine if a new priority number PNEW 1710 of a new policy statement 1708 is the same, higher, or, alternatively, lower than any of the priority numbers already stored in memory 608. The inequality circuits also update the priorities of the existing policy statements in order to insert the new policy statement. The new policy statement and its associated priority number do not need to be physically inserted between the existing entries; rather, they can be loaded into any desirable address (e.g., the next free address) in the CAM array and priority memory, and the priority numbers of the existing policy statements updated accordingly. Similarly, when a policy statement is removed (i.e., invalidated or overwritten) from the CAM array, the inequality circuits may update the priority numbers of the previously stored policy statements. These updating functions can be performed by the inequality circuits without the need to physically reorder the policy statements in the CAM array, or to physically reorder the priority numbers in the priority index table. This can reduce the hardware and/or software needed for table management of the CAM array, and can increase the performance of a router incorporating the DSP.

For an alternative embodiment, the inequality circuits do not update the existing, stored priority numbers during or after an insert or delete operation. For this embodiment, gaps may be left in the priority numbers after an insert or

delete function. These gaps may or may not be filled with subsequent priority numbers.

As shown in Figure 17, each location 608_0 - 608_{N-1} in the priority memory has a corresponding inequality circuit 1704_0 - 1704_{N-1} that compares the new priority number PNEW 1710 with the priority number P_0 - P_Z stored at that location. If PNEW is the same or a higher priority (e.g., is greater than or equal to, or, alternatively less than or equal to) than an existing priority number, an inequality circuit will update the existing priority number. The new policy statement 1708, PNEW, and the associated route information RI_{NEW} , can then be written into CAM array 604, priority memory 608, and route memory 614, respectively, using conventional write circuits, counters, address decoders, etc. (not shown).

Inequality circuits 1704_0 - 1704_{N-1} may be part of priority index table 606 as shown in Figure 17. Alternatively, inequality circuits 1704_0 - 1704_{N-1} may be separate from priority index table 606. The inequality circuits may be any inequality circuit that determines whether one number is greater than or equal to, greater than, less than or equal to, or less than another number. For an alternative embodiment, inequality circuits 1704_0 - 1704_{N-1} may be one circuit that consecutively, or in parallel, compares PNEW with each of the priority numbers P_0 - P_Z .

Figure 18 shows one example of adding new policy statement 1708 to CAM array 604 at the next free address of CAM array 604 (location 604_3) and updating the existing priority numbers stored in memory 608. New policy statement 1708 is assigned priority number 1, and has route information RI_3 . CAM array 604 already includes policy statements 201, 202, and 203 stored at locations 604_0 , 604_1 , and 604_2 , respectively. The corresponding priority numbers 2, 1, and 0 are stored at locations 608_0 , 608_1 , and 608_2 , respectively, in priority

memory 608. Additionally, the corresponding route information RI_0 , RI_1 , and RI_2 , are stored at locations 614_0 , 614_1 , and 614_2 , respectively, of route memory 614. For this embodiment, the policy statements are assigned priority numbers in ascending priority order such that policy statement 203 is identified as the

5 highest priority policy statement by being assigned priority number 0; policy statement 201 is identified as the lowest priority policy statement by being assigned priority number 2; and, policy statement 202 is identified as having a priority greater than that of policy statement 201, but less than that of policy statement 203, by being assigned priority number 1.

10 Inequality circuits 1704_0 - 1704_2 compare PNEW with the exiting priority numbers as follows. The comparisons are made in parallel by each inequality circuit, however, they may also be made consecutively or in groups. Inequality circuit 1704_2 compares its corresponding priority number 0 with the new priority number 1 and determines that 0 is a higher priority than 1 (i.e., that 0 is

15 numerically less than 1). The priority number for policy statement 203 is not updated. Inequality circuit 1704_1 compares its corresponding priority number 1 with the new priority number 1 and determines that they have the same priority. The priority number for policy statement 202 is incremented (updated) by one to the number 2. Inequality circuit 1704_0 compares its corresponding priority

20 number 2 with the new priority number 1 and determines that 1 is a higher priority than 2 (i.e., that 1 is numerically less than 2). The priority number for policy statement 201 is then incremented to 3. The modified table entries after the insertion of the new policy statement and its associated priority number are shown in Figure 19.

25 Figure 20 shows another example of adding new policy statement 1708 to CAM array 604 at the next free address of CAM array 604 (location 604_3). For this embodiment, policy statements 201, 202, and 203 have priority numbers 14,

15, and 16, respectively, and the new policy statement has priority number 15. The policy statements are assigned in descending priority order such that policy statement 203 is identified as the highest priority policy statement by being assigned priority number 16; policy statement 201 is identified as the lowest priority policy statement by being assigned priority number 14; and, policy statement 202 is identified as having a priority greater than that of policy statement 201, but less than that of policy statement 203, by being assigned priority number 15.

Inequality circuit 1704₂ compares its corresponding priority number 16 with the new priority number 15 and determines that 16 is a higher priority than 15 (i.e., that 16 is numerically greater than 15). The priority number for policy statement 203 is not updated. Inequality circuit 1704₁ compares its corresponding priority number 15 with the new priority number 15 and determines that they have the same priority. The priority number for policy statement 202 is decremented (updated) by one to the number 14. Inequality circuit 1704₀ compares its corresponding priority number 14 with the new priority number 15 and determines that 15 is a higher priority than 14 (i.e., that 14 is numerically less than 15). The priority number for policy statement 201 is then decremented to 13. The modified table entries after the insertion of the new policy statement and its associated priority number is shown in Figure 21.

As indicated above, the priority numbers can be incremented, decremented, or left unchanged after PNEW is compared with the priority number. Figure 22 shows priority memory 2208 that is one embodiment of priority memory 608. Priority memory 2208 includes counters 2208₀-2208_{N-1}. Each counter corresponds to priority memory location 608₀-608_{N-1}, respectively, and stores the priority number for that location. The counters can be loaded (i.e., written to) using write circuits, read from using read circuits, or reset as

generally known in the art. Additionally, each counter 2208₀-2208_{N-1} provides its stored priority number P₀-P_Z to a corresponding inequality circuit 1704₀-1704_{N-1} for comparison with PNEW. When PNEW has a higher priority than the stored priority number for a given location, then the inequality circuit will assert its count signal COUNT₀-COUNT_{N-1} such that the corresponding counter is incremented (or, alternatively, decremented). The counter will be incremented if the UP signal is asserted, or will be decremented if the DOWN signal is asserted. The UP and DOWN signals may be generated by an instruction decoder (not shown) on DSP 602 that receives and decodes an insert (write) instruction or a delete (invalidate) instruction to the DSP. For an alternative embodiment, the inequality circuits may each output UP and DOWN signals to their corresponding counters directly in response to one or more signals from the instruction decoder. The counters may be updated synchronously or asynchronously.

Inequality circuits 1704₀-1704_{N-1} may be disposed external to priority memory 2208, or they may be integrated into the priority memory itself. Figure 23 shows one embodiment of integrating an inequality circuit with a counter in the priority memory. For this embodiment, inequality circuit 2304₀ is one embodiment of inequality circuit 1704₀, and may also be used for any of inequality circuits 1704₁-1704_{N-1}. Counter 2306₀ is one embodiment of counter 2208₀, and may also be used for any of counters 2208₁-2208_{N-1}.

As shown in Figure 23, memory storage elements 2314_{n-1,0}-2314_{0,0} have associated counter logic circuits 2316_{n-1,0}-2316_{0,0} to form a conventional counter. Each memory storage element can be a volatile or non-volatile RAM or ROM storage element such as memory storage elements 1102 of Figure 11. The memory storage elements each store one bit of a priority number such that memory storage element 2314_{n-1,0} stores the most significant priority number bit,

2314_{n-2,0} stores the next most significant priority number bit, and 2314_{0,0} stores the least significant priority number bit. Logic circuits 2316_{n-1,0}-2316_{0,0} receive the UP, DOWN, and COUNT₀ signals that indicate when the counter should increment, decrement, or leave unchanged the priority number stored in the counter. The counter can be configured to form any type of counter including the up/down counter shown. Additionally, load and reset signals can be provided to the logic circuits to write or reset the counter. Bit lines, word lines, and read/write circuits generally known in the art may also be included to access the priority number bits stored in the memory storage elements.

Inequality circuit 2304₀ includes compare circuits 2310_{n-1,0}-2310_{0,0}. Each compare circuit can be any digital or analog compare circuits. Each compare circuit 2310_{n-1,0}-2310_{0,0} compares a priority number bit, stored in a corresponding memory storage element 2314_{n-1,0}-2314_{0,0}, with one of the corresponding bits PNEW_{n-1}-PNEW₀ provided on signal lines 2308_{n-1}-2308₀. Additionally, each compare circuit monitors the comparison result of the more significant priority number bits with PNEW via the logical states of control line segments 2312_{n-2,0}-2312_{0,0}. The first control line segment 2312_{n-1,0} may be pre-charged by pre-charge circuit 2302 to indicate that it is the first control line segment. Alternatively, control line segment 2312_{n-1,0} may be initially discharged to indicate that it is the first control line segment.

For this embodiment, the logical state of COUNT₀ on signal line 2318 indicates the comparison result between PNEW and the priority number stored in counter 2306₀. COUNT₀ is initially pre-charged by a pre-charge circuit 2302 to a high logic state. If COUNT₀ remains pre-charged after each compare circuit performs its comparison function, then the priority number in counter 2306₀ is updated. Alternatively, if COUNT₀ is discharged by one of the compare circuits, then the priority number in counter 2306₀ is updated. For other embodiments,

COUNT₀ may be initially discharged to a logic low state, and counter 2306 updated if COUNT₀ remains discharged, or, alternatively, if COUNT₀ is charged to a logic one state by one of the compare circuits. For yet another embodiment, signal line 2318 may be segmented like control line segments 2312 between
5 compare circuits, and the segmented signal line 2318 output from compare circuit 2310_{0,0} provided to counter logic circuit 2316_{0,0} as COUNT₀.

The operation of the embodiment of Figure 23 can be illustrated with the example shown in Figure 24. In this example, a 3x3 matrix of rows and columns stores three priority numbers. For other embodiments, any numbers of rows and
10 columns can be used. Row 0 stores priority number 010 having the decimal equivalent of the number 2, row 1 stores priority number 001 having the decimal equivalent of the number 1, and row 2 stores priority number 000 having the decimal equivalent of the number 0. For this example, the priority numbers are assigned in ascending priority order such that 000 is the most significant priority
15 number, 001 is the next most significant priority number, and 010 is the least significant priority number.

When a new policy statement having PNEW 001 is to be loaded into CAM array 604, inequality circuits 2304₂-2304₀ compare PNEW with their corresponding priority numbers stored in counters 2306₂-2306₀, respectively.

20 Inequality circuit 2304₂ determines that 000 is a higher priority than PNEW; inequality circuit 2304₁ determines that 001 is equal to PNEW; and inequality circuit 2304₀ determines that 010 is a lower priority than PNEW. Inequality circuit 2304₂ leaves its priority number unchanged. Inequality circuits 2304₁ and 2304₀ cause counters 2306₁ and 2306₀ to increment by one their existing priority
25 numbers such that they do not have the same priority number as PNEW. The new policy statement and PNEW can then be loaded into CAM array 604 and priority memory 608, respectively.

Inequality circuits 2304₁ and 2304₀ cause their corresponding counters 2306₁ and 2306₀ to be incremented as follows. The comparisons of the most significant bit PNEW₂ with the most significant bits stored in counters 2306₂-2306₀ are resolved first. When an individual compare circuit determines that its
 5 corresponding memory storage element stores the same logic state as the PNEW bit, the compare circuit leaves the COUNT signal line unaffected, and drives the next control line segment to the same logical state of the previous control line segment. Since each of memory storage elements 2314_{2,2}, 2314_{2,1}, and 2314_{2,0} store a logic zero and PNEW₂ is a logic zero, COUNT₂, COUNT₁, and COUNT₀
 10 remain unaffected. Additionally, control signals on control line segments 2312_{1,2}, 2312_{1,1}, and 2312_{1,0} are asserted to high logic states. The inequality circuits cannot yet determine, based on the most significant bits alone, whether PNEW is greater than or equal to any of the stored priority numbers.

The comparison of the next most significant bit PNEW₁ with priority
 15 number bits stored in memory elements 2314_{1,0}, 2314_{1,1} and 2314_{1,2} is then resolved. Since memory elements 2314_{1,2} and 2314_{1,1} store the same logic states as PNEW₁, COUNT₂ and COUNT₁ remain unaffected, and control signals on control line segments 2312_{0,2} and 2310_{0,1} are asserted to high logic states. Compare circuit 2310_{1,0}, however, determines that the priority number stored in
 20 counter 2306₀ is greater than PNEW because: (1) memory element 2314_{1,0} stores a logic one; (2) the control signal is high on control line segment 2312_{1,0}; and (3) PNEW₁ is a logic zero on signal line 2308₁. When compare circuit 2310_{1,0} makes this determination, it leaves COUNT₀ unaffected so that it will remain pre-charged to update the priority number stored in counter 2306₀. Additionally,
 25 compare circuit 2310_{1,0} asserts the control signal on control line segment 2312_{0,0} to a low logic state. When the control signal on control line segment 2310_{0,0} is low, the control signals on all subsequent control line segments are also asserted to

low logic states. When the control signals are low, all lower priority compare circuits are rendered incapable of discharging $COUNT_0$ regardless of what is stored in their corresponding memory elements, and regardless of the logical states of the subsequent PNEW bits. For example, $2310_{0,0}$ will be rendered

5 incapable of discharging $COUNT_0$ even though $PNEW_0$, a logic one, is greater than the logic zero stored in memory element $2314_{0,0}$.

Lastly, the comparison of the least significant bit $PNEW_0$ with priority number bits stored in memory elements $2314_{0,0}$, $2314_{0,1}$ and $2314_{0,2}$ is resolved. As indicated above, inequality circuit 2304_0 has already determined that 010 is

10 greater than 001 and thus $COUNT_0$ remains pre-charged to update the priority number stored in counter 2306_0 . Since the control signal on control line segment $2312_{0,1}$ is high, compare circuit $2310_{0,1}$ compares the logic one stored in memory element $2314_{0,1}$ with the logic one of $PNEW_0$ and determines that PNEW is the same number as that stored in counter 2306_1 . In response, $COUNT_1$ remains in a

15 pre-charged state to update the priority number stored in counter 2306_1 . Compare circuit $2310_{0,2}$, however, determines that $PNEW_0$ is greater than the logic zero stored in memory element $2314_{0,2}$. Since the control signal is high on control line segment $2312_{0,2}$, compare circuit $2310_{0,2}$ discharges $COUNT_2$ such that the priority number 000 stored in counter 2306_2 will not be updated. After a

20 sufficient amount of time has passed such that all of the bit comparisons have had time to resolve, the UP signal can be asserted such that the priority numbers in counters 2306_1 and 2306_0 are incremented. For one embodiment, the COUNT signals can be latched prior to providing the UP signal.

Any compare circuits may be used for compare circuits 2310 to implement

25 the process illustrated above. Table 3 shows one example of a truth table for implementing each compare circuit 2310, where X (column) and Y (row) are any integers. Other truth tables may be used (and corresponding logic generated

accordingly) including those that logically complement one or more of the signals indicated in Table 3.

| PNEW _x | 2314 _{xY} | 2312 _{xY} | 2312 _{x-1,Y} | COUNT |
|-------------------|--------------------|--------------------|-----------------------|-------|
| 0 | 0 | 0 | 0 | COUNT |
| 0 | 0 | 1 | 1 | COUNT |
| 0 | 1 | 0 | 0 | COUNT |
| 0 | 1 | 1 | 0 | COUNT |
| 1 | 0 | 0 | 0 | COUNT |
| 1 | 0 | 1 | 1 | 0 |
| 1 | 1 | 0 | 0 | COUNT |
| 1 | 1 | 1 | 1 | COUNT |

Table 3

Any logic or circuitry may be used to implement the truth table of Table 3. Figure 25A shows compare circuit 2502 that is one embodiment of compare circuit 2310_{n-1,0} for implementing the truth table of Table 3. The embodiment of Figure 25A may be used to implement all of the compare circuits 2310.

Compare circuit 2502 includes inverter 2510, NAND gate 2508, AND gate 2506, and n-channel transistors 2504, 2503, and 2507. NAND gate 2508 has one input coupled to signal line 2308_{n-1} via inverter 2510, and another input coupled to receive the data (D) stored in memory element 2314_{n-1,0}. AND gate 2506 has one input coupled to the output of NAND gate 2508, another input coupled to control line segment 2312_{n-1,0}, and an output coupled to control line segment 2312_{n-2,0}. Transistor 2504 has its drain coupled to the signal line 2318, its gate coupled to the output of AND gate 2506, and its source coupled to the drain of transistor 2503. Transistor 2503 also has its gate controlled by signal line 2308_{n-1}, and its source coupled to the drain of transistor 2507. Transistor 2507 also has its

gate coupled to receive /D from memory 2314_{n-1,0}, and its source coupled to ground.

Figure 25B shows compare circuit 2522 that is another embodiment of compare circuit 2310_{n-1,0} for implementing the truth table of Table 3. In this embodiment, signal line 2318 is segmented into separate control line segments. The embodiment of Figure 25B may be used to implement all of the compare circuits 2310.

Compare circuit 2522 includes inverter 2512, NAND gates 2514 and 2516, and AND gates 2518 and 2520. NAND gate 2514 has one input coupled to signal line 2308_{n-1} via inverter 2512, and another input coupled to D of memory 2314_{n-1,0}. AND gate 2520 has one input coupled to control line segment 2312_{n-1,0}, another input coupled to the output of NAND gate 2514, and an output coupled to control line segment 2312_{n-2,0}. NAND gate 2516 has a first input coupled to control line segment 2312_{n-1,0}, a second input coupled to /D of memory 2314_{n-1,0}, and a third input coupled to signal line 2308_{n-1}. AND gate 2518 has a first input coupled to count line segment 2318_{n-1,0}, a second input coupled to the output of NAND gate 2516, and an output coupled to count line segment 2318_{n-2,0}.

For the example described above with respect to Figure 24, the most significant number is the lowest number such that 000 is the most significant number between 000, 001, and 010. For another embodiment, the priority numbers are assigned in descending priority order such that the most significant priority number is the priority number with the highest numerical value, and the least significant priority number is the priority number with the lowest numerical value. Figure 26 shows such an example in which the 3x3 matrix of Figure 24 includes the binary priority number 111 (decimal equivalent of 7) stored in counter 2306₂, priority number 110 (decimal equivalent of 6) stored in counter 2306₁, and priority number 101 (decimal equivalent of 5) stored in

counter 2306₀. Numbers 7, 6, and 5 are stored in rows 2, 1, and 0, respectively. In this example, 111 is the most significant priority number, 110 is the next most significant priority number, and 101 is the least significant priority number.

When a new policy statement having PNEW 110 is to be loaded into CAM array 604, the logical complement of PNEW (i.e., 001) is provided to signal lines 2308₂-2308₀. Inequality circuits 2304₂-2304₀ compare PNEW with their corresponding priority numbers stored in counters 2306₂-2306₀, respectively. Inequality circuit 2304₂ determines that 111 is a higher priority than PNEW; inequality circuit 2304₁ determines that 110 is equal to PNEW; and inequality circuit 2304₀ determines that 101 is a lower priority than PNEW. Inequality circuit 2304₂ leaves its priority number unchanged. Inequality circuits 2304_{0,1} and 2304_{0,0} each cause their corresponding counters 2306₁ and 2306₀ to decrement by one their existing priority numbers such that they do not have the same priority number as PNEW. The new policy statement and PNEW can then be loaded into CAM array 604 and priority memory 604, respectively.

Inequality circuits 2304₁ and 2304₀ cause their corresponding counters 2306₁ and 2306₀ to be decremented as follows. The comparisons of the most significant bit PNEW₂ with the most significant bits stored in counters 2306₂-2306₀ are resolved first. When an individual compare circuit determines that its corresponding memory storage element stores the same logic state as the PNEW bit, the compare circuit leaves the COUNT signal line unaffected, and drives the next control line segment to the same logical state of the previous control line segment. Since each of memory storage elements 2314_{2,2}, 2314_{2,1}, and 2314_{2,0} store a logic one and PNEW₂ is a logic one, COUNT₂, COUNT₁, and COUNT₀ remain unaffected. Additionally, control signals on control line segments 2312_{1,2}, 2312_{1,1}, and 2312_{1,0} are asserted to high logic states. The inequality circuits cannot yet

determine, based on the most significant bits alone, whether PNEW is less than or equal to any of the stored priority numbers.

The comparison of the next most significant bit $PNEW_1$ with priority number bits stored in memory elements $2314_{1,0}$, $2314_{1,1}$ and $2314_{1,2}$ is then

5 resolved. Since memory elements $2314_{1,2}$ and $2314_{1,1}$ store the same logic states as $PNEW_1$, $COUNT_2$ and $COUNT_1$ remain unaffected, and control signals on control line segments $2312_{0,2}$ and $2310_{0,1}$ are asserted to high logic states.

Compare circuit $2310_{1,0}$, however, determines that the priority number stored in counter 2306_0 is less than PNEW because: (1) memory element $2314_{1,0}$ stores a

10 logic zero; (2) the control signal is high on control line segment $2312_{1,0}$; and (3) $PNEW_1$ is a logic one. When compare circuit $2310_{1,0}$ makes this determination, it leaves $COUNT_0$ unaffected so that it will remain pre-charged to update the priority number stored in counter 2306_0 . Additionally, compare circuit $2310_{1,0}$ asserts the control signal on control line segment $2312_{0,0}$ to a low logic state.

15 When the control signal on control line segment $2310_{0,0}$ is low, the control signals on all subsequent control line segments are also asserted to low logic states.

When the control signals are low, all lower priority compare circuits are rendered incapable of discharging $COUNT_0$ regardless of what is stored in their corresponding memory elements, and regardless of the logical states of the

20 subsequent PNEW bits. For example, $2310_{0,0}$ will be rendered incapable of discharging $COUNT_0$ even though $PNEW_0$, a logic zero, is less than the logic one stored in memory element $2314_{0,0}$.

Lastly, the comparison of the least significant bit $PNEW_0$ with priority number bits stored in memory elements $2314_{0,2}$, $2314_{0,1}$ and $2314_{0,0}$ is resolved. As

25 indicated above, inequality circuit 2304_0 has already determined that 101 is less than 110 and thus $COUNT_0$ remains pre-charged to update the priority number stored in counter 2306_0 . Since the control signal on control line segment $2312_{0,1}$ is

high, compare circuit 2310_{0,1} compares the logic zero stored in memory element 2314_{0,1} with the logic one on line 2308₀ and determines that PNEW is the same number as that stored in counter 2306₁. In response, COUNT₁ remains in a pre-charged state to update the priority number stored in counter 2306₁. Compare circuit 2310_{0,2}, however, determines that PNEW₀ is less than the logical one stored in memory element 2314_{0,2}. Since the control signal is high on control line segment 2312_{0,2}, compare circuit 2310_{0,2} discharges COUNT₂ such that the priority number 111 stored in counter 2306₂ will not be updated. After a sufficient amount of time has passed such that all of the bit comparisons have had time to resolve, the DOWN signal can be asserted such that the priority numbers in counters 2306₁ and 2306₀ are decremented. For one embodiment, the COUNT signals can be latched prior to providing the DOWN signal.

Any compare circuits may be used for compare circuits 2310 to implement the process illustrated above. Table 4 shows one example of a truth table for implementing each compare circuit 2310, where X (column) and Y (row) are any integers. Other truth tables may be used (and corresponding logic generated accordingly) including those that logically complement one or more of the signals indicated in Table 4.

20

25

| PNEW _x | 2314 _{x,y} | 2312 _{x,y} | 2312 _{x-1,y} | COUNT |
|-------------------|---------------------|---------------------|-----------------------|-------|
| 0 | 0 | 0 | 0 | COUNT |
| 0 | 0 | 1 | 1 | COUNT |
| 0 | 1 | 0 | 0 | COUNT |
| 0 | 1 | 1 | 1 | 0 |
| 1 | 0 | 0 | 0 | COUNT |
| 1 | 0 | 1 | 0 | COUNT |
| 1 | 1 | 0 | 0 | COUNT |
| 1 | 1 | 1 | 1 | COUNT |

Table 4

Any logic or circuitry may be used to implement the truth table of Table 4. Figure 27A shows compare circuit 2702 that is one embodiment of compare circuit 2310_{n-1,0} for implementing the truth table of Table 4. The embodiment of Figure 27A may be used to implement all of the compare circuits 2310.

Compare circuit 2702 includes inverter 2510, NAND gate 2508, AND gate 2506, and n-channel transistors 2504, 2503, and 2507 as in Figure 25A. NAND gate 2508 has one input coupled to signal line 2308_{n-1}, and another input coupled to receive the /D stored in memory element 2314_{n-1,0}. AND gate has one input coupled to the output of NAND gate 2508, another input coupled to control line segment 2312_{n-1,0}, and an output coupled to control line segment 2312_{n-2,0}.

Transistor 2504 has its drain coupled to the signal line 2318, its gate coupled to the output of AND gate 2506, and its source coupled to the drain of transistor 2503. Transistor 2503 also has its gate controlled by signal line 2308_{n-1} via inverter 2510, and its source coupled to the drain of transistor 2507. Transistor 2507 also has its gate coupled to receive D from memory 2314_{n-1,0}, and its source coupled to ground.

Figure 27B shows compare circuit 2704 that is another embodiment of compare circuit 2310_{n-1,0} for implementing the truth table of Table 4. In this embodiment, signal line 2318 is segmented into separate control line segments. The embodiment of Figure 27B may be used to implement all of the compare
5 circuits 2310.

Compare circuit 2704 includes inverter 2708, NAND gate 2710, and AND gates 2706 and 2712. NAND gate 2710 has a first input coupled to signal line 2308_{n-1}, a second input coupled to D of memory 2314_{n-1,0}, and a third input coupled to control line segment 2312_{n-1,0} via inverter 2708. AND gate 2706 has
10 one input coupled to count line segment 2318_{n-1,0}, another input coupled to the output of NAND gate 2710, and an output coupled to count line segment 2318_{n-2,0}. AND gate 2712 has a first input coupled to control line segment 2312_{n-1,0}, a second input coupled to /D of memory 2314_{n-1,0}, a third input coupled to signal line 2308_{n-1}, and an output coupled to control line segment 2312_{n-2,0}.

The embodiments of Figures 17-27 show that a new policy statement having a priority greater than or equal to an existing policy statement, can be added to the DSP without having to physically reorder or reload the CAM array storing the policy statements, or the memory storing the priority numbers. Instead, inequality circuits update the appropriate existing priority numbers.
15 Since the updating function can be performed without the need to physically reorder or reload the policy statements in the CAM array, or to physically reorder or reload the priority numbers in the priority index table, this can reduce the hardware and/or software needed for table management of the CAM array, and can increase the performance of a router incorporating the DSP.

Policy statements may also be deleted from CAM array 604 using priority index table 606 of Figure 17. When a policy statement is deleted, the corresponding priority number in priority memory 608 is also deleted. The
25

policy statements and priority numbers can be deleted by setting one or more valid bits to an appropriate state for the row of CAM array 604 that stores the policy statement to be deleted. The valid bit(s) may be stored in CAM array 604, priority memory 608, or in each of the memory arrays.

- 5 When a priority number is deleted, the remaining priority numbers in priority memory 608 can be left unchanged. This may leave gaps in the priority numbers stored in the priority memory. These gaps may be filled in by new priority numbers associated with new policy statements to be added to the DSP, or they may remain as unused gaps. For another embodiment, remaining
- 10 priority numbers in the priority memory can be updated by the inequality and counter circuits. If the priority numbers are assigned in ascending priority order, and one of the priority numbers is deleted or invalidated, then any other priority numbers that are greater than the deleted number can be decremented by the inequality circuits. For example, if priority memory 608 stores priority
- 15 numbers 0, 1, and 2, and priority number 1 is invalidated, then priority number 2 can be decremented to 1. Similarly, if the priority numbers are assigned in descending priority order, and one of the priority numbers is deleted or invalidated, then any other priority numbers that are less than the deleted
- 20 number can be incremented by the inequality and counter circuits. For example, if priority memory 608 stores priority numbers 0, 1, and 2, and priority number 1 is invalidated, then priority number 0 can be incremented to 1. Updating the priority numbers can help to avoid unused gaps in priority memory 608.

- 25 Figure 28 shows one embodiment of a process 2800 for deleting or invalidating policy statements and their associated priority numbers. At step 2802, the priority number to be deleted is compared with the entries in the priority memory. For this embodiment, priority memory 608 is a CAM, or has one or more compare circuits external to the memory to perform the comparison.

For an alternative embodiment, step 2802 may compare the policy statement with the existing entries in CAM array 604. For this embodiment, priority memory 608 may not be a CAM.

If there is no match at step 2802, the process stops at step 2804. The process could also stop if there was more than one match, and an error flag such as a multiple match flag may be asserted. Alternatively, the process may continue to delete all identified matches. If there is a match, both the policy statement and the priority number are invalidated or deleted at step 2806. As described above, this may be accomplished by setting a valid bit(s) to an appropriate state for that row in priority memory 608. If the valid bit(s) is located in CAM array 604, and is inaccessible by priority memory 608, then priority logic 610 can then identify the location of the matching entry (as described above), and provide the indication to IAD_0 - IAD_{N-1} . These signal lines can then be provided through encoder 612 back to CAM array 604 (e.g., through a decoder) to address the appropriate location and invalidate the policy statement. Alternatively, the signal lines IAD_0 - IAD_{N-1} can be provided directly to CAM array 604 to access the appropriate location for invalidation.

The process may stop after step 2806 and potentially leave gaps in the priority numbers. Alternatively, the process may proceed to step 2808 to update the remaining entries in the priority memory. When the process continues, the deleted priority number is provided to the inequality circuits 1704_0 - 1704_{N-1} , at step 2808, to determine which entries in the priority memory have a lower priority than the deleted priority number. The inequality and counter circuits may then update the identified numbers, at step 2810, as previously described.

The delete functions illustrate one example of the DSP 602 utilizing both the inequality circuits and the priority logic to operate on the priority numbers stored in priority memory 608. Figure 29 shows one example of combining an

inequality circuit with a priority logic circuit into a single array such that each circuit operates on the priority numbers on a bit-by-bit basis. The embodiment of Figure 29 combines row 1201₀ of priority logic elements from Figure 12A with inequality circuit 2304₀ and counter 2306₀ of Figure 23 to form row zero in

5 priority index table 606 of Figure 17. For an alternative embodiment, row 1201₀ from Figure 12B may also be used. The priority logic elements and inequality circuits share memory elements 2314_{0,0}-2314_{n-1,0}. Memory elements 2314_{0,0}-2314_{n-1,0} also correspond to memory elements 1102_{0,0}-1102_{n-1,0}. Conventional bit lines, word lines, read and write circuits, and optional comparand data lines and
10 match lines (i.e., when memory elements 2314_{0,0}-2314_{n-1,0} are CAM cells) are not shown so as not to obscure the teachings of the figure.

Figure 30 shows another embodiment of combining the row 1201₀ of priority logic elements from Figure 12A with the inequality circuit 2304₀ of Figure 23 to form row zero in priority index table 1706 of Figure 17. In this
15 embodiment, signal lines 1208_{n-1}-1208₀ and 2308_{n-1}-2308₀ are replaced with a single set of signal lines 2906_{n-1}-2906₀ that can be used to determine the most significant priority number PNUM, or used to carry PNEW. Also, for this embodiment, match line segments 1210_{n-1,0}-1210_{0,0} serve a dual function to operate as match line segments when determining the most significant priority
20 number, and as control line segments (i.e., 2312_{n-1,0}-2312_{0,0} of Figure 23) to generate COUNT₀ when determining when a new priority number has a priority that is greater than or equal to the priority number stored in counter 2306. Thus, compare circuit 2310_{n-1,0} is coupled to match line segment 1210_{n-2,0}, compare circuit 2310_{n-2,0} is coupled to match line segment 1210_{n-3,0}, and so forth, and
25 COUNT₀ is provided on signal line IAD₀. For an alternative embodiment, compare circuits 2310 may provide control signals to isolation circuits 1204 instead of directly controlling the match line segments.

A control signal BCT indicates whether the priority index table uses the row of priority logic elements 1201₀ to determine the most significant priority number in priority memory 608, or whether the priority index table uses the inequality circuit 2304₀ to determine if a new priority number has a priority that is greater than or equal to the priority number stored in counter 2306. BCT may be provided by an instruction decoder in DSP 602 that interprets and decodes instructions provided to the DSP. When BCT is a low logic state, the priority logic elements perform their priority function; and when BCT is a high logic state, the inequality circuits perform their inequality function. Thus, when BCT is low, compare circuit 1206_{n-1,0} is enabled, compare circuit 2310_{n-1,0} is disabled, and logic 2902 provides the state of ML₀ to match line segment 1210_{n-1,0}. When BCT is high, compare circuit 1206_{n-1,0} is disabled, compare circuit 2310_{n-1,0} is enabled, and logic 2902 drives match line segment 1210_{n-1,0} to a high logic state such that match line segments 1210_{n-1,0}-1210_{0,0} function as control line segments 2312_{n-1,0}-2312_{0,0} of Figure 23. Logic 2902 may be any logic that performs this function.

For another embodiment of Figure 30, row 1201₀ of priority logic elements from Figure 12B may be combined with the inequality circuit 2304₀ of Figure 23 to form row zero in priority index table 1706 of Figure 17.

The logic and circuitry for implementing the priority logic elements and the inequality circuits illustrated in Figures 29 and 30 may be those that implement the truth tables shown in Tables 1, 2, 3, and 4 above. For example, the logic and circuits illustrated in Figures 14, 15, 25A, 25B, 27A, and 27B may be used.

For one embodiment, the priority numbers are assigned in ascending priority order in priority memory 608. For this embodiment, logic and circuitry that implements the truth table shown in Table 1 (e.g., the logic and circuitry of

Figures 14 or 15) may be used for the priority logic elements to determine the lowest priority number as the most significant priority number. Also for this embodiment, the logic and circuitry shown in the truth table shown in Table 3 (e.g., the logic and circuitry of Figures 25A or 25B) may be used for the

5 inequality circuits to determine if a new priority number has a priority that is greater than or equal to the priority number stored in counter 2306.

Figure 31 shows one example of combining the logic and circuitry of Figures 14 and 25 to implement the embodiment of Figure 30 when the priority numbers are assigned in ascending priority order in priority memory 608. For

10 this example, OR gate 3102 is one embodiment of logic 2902. OR gate 3102 has one input coupled to ML_0 , and another input coupled to BCT. BCT is also coupled to the gate of n-channel transistor 3106 and one input of OR gate 3104 via inverter 3108. Transistor 3106 has its drain coupled to signal line 2906_{n-1} , and its source coupled to the drain of transistor 1406. OR gate 3104 has another input

15 coupled to the output of AND gate 2506, and has its output coupled to the gate of transistor 2504. Transistor 2504 is coupled between line segment $1210_{n-2,0}$ and node 3112. Transistor pairs 1410/1412 and 2503/2507 are each coupled in series between node 3112 and ground.

For another embodiment, the priority numbers are assigned in

20 descending priority order in priority memory 608. For this embodiment, logic and circuitry that implements the truth table shown in Table 2 (e.g., the logic and circuitry of Figures 14 or 15 modified as indicated above) may be used for the priority logic elements to determine the lowest priority number as the most significant priority number. Also for this embodiment, the logic and circuitry

25 shown in the truth table shown in Table 4 (e.g., the logic and circuitry of Figures 27A or 27B) may be used for the inequality circuits to determine if a new priority number has a priority that is less than or equal to the priority number stored in

counter 2306. Figure 32 shows one example of combining the logic and circuitry of Figures 14 and 27 to implement the embodiment of Figure 30 when the priority numbers are assigned in descending priority order in priority memory 608.

5 The embodiments of Figures 17-32 include additional inequality circuits to compare PNEW with the priority numbers already stored in the priority index table. For another embodiment, priority logic 610 itself can be used to determine whether to update existing priority numbers stored in the priority index table in response to an insert or delete instruction. As with the previously described

10 embodiments above, the new policy statement and PNEW do not need to be physically inserted between the existing entries; rather, they can be loaded into any desirable address (e.g., the next free address) in the CAM array and priority memory, and the priority numbers of the existing policy statements updated accordingly. Similarly, when a policy statement is removed (i.e., invalidated or

15 overwritten) from the CAM array, the priority logic may update the priority numbers of the previously stored policy statements. These updating functions can be performed by the priority logic without the need to physically reorder the policy statements in the CAM array, or to physically reorder the priority numbers in the priority index table. This can reduce the hardware and/or

20 software needed for table management of the CAM array, and can increase the performance of a router incorporating the DSP.

 When the priority numbers are assigned in ascending priority order and PNEW is to be added to any free location in memory 608, priority logic 610 can determine whether the priority numbers already stored in memory 608 should

25 be incremented as follows. PNEW is decremented by one (or another value) and then compared with the existing priority numbers stored in memory 608 by priority logic 610. For this embodiment, the priority numbers in memory 608 are

not compared with each other, but with the decremented PNEW. Since priority logic 610 determines which compared number is the lowest numerical value, it can also determine the logical converse; that is, it can determine which of the stored priority numbers are greater than the decremented PNEW. Each priority number that is greater than the decremented PNEW is then incremented by the priority logic. PNEW is then added to memory 608.

Figure 33 shows one example of inserting policy statement 1708 and PNEW into memory 608 having priority numbers 0, 1, and 2 assigned in ascending priority order. This is the same example shown in Figures 18 and 19, except that priority logic 610 is used to control updating of the priority numbers stored in memory 608. PNEW, having a value of 1, is decremented by 1 by decrement circuit 3304 to form PSUB having a value of zero. PSUB is then provided to priority table 606 via multiplexer 3302. Decrement circuit 3304 may be any subtraction circuit that decrements PNEW by one or another value. Multiplexer 3302 provides PSUB to priority index table 606 in response to control signal INST, which indicates that an insert function is to take place. INST may be generated by an instruction decoder (not shown) that receives an insert (write) instruction.

PSUB is compared with the stored priority numbers by priority logic 610. In response to control signal BCT, logic 3306 asserts all match lines ML_0 - ML_{N-1} to a particular logic state to indicate that all stored policy statements match policy statement 1708 (even though no actual comparison takes place) such that the priority logic will compare all of the priority numbers in memory 608 with PSUB. Logic 3306 may be any logic circuit including, for example, AND, OR, XOR, or inverter circuits. BCT is a control signal that indicates an insert or delete function is being performed by DSP 602. BCT may be generated by an instruction decoder (not shown). BCT may also be provided to priority logic 610

to indicate that priority logic 610 should compare PSUB with each entry in memory 608 rather than comparing each entry in memory 608 with each other as in the classification function of Figure 7.

Priority logic 610 determines that stored priority number 0 is equal to PSUB, and leaves this stored number unchanged. Priority logic 610 also determines that stored priority numbers 1 and 2 are larger than PSUB, and increments these numbers to 2 and 3, respectively. Policy statement 1708 and PNEW are then loaded into a free address in array 604 and memory 608, respectively. The modified table entries after the insertion are shown in Figure 34.

Note that if the priority number 0 is to be inserted, then decrement circuit 3304 would generate a negative number. For one embodiment, only priority numbers greater than 0 may be used. For another embodiment, offset circuitry may be included in DSP 602 to add an offset value to PNEW before it is provided to decrement circuit 3304 and/or multiplexer 3302. The offset circuitry may also be used to subtract the offset value when reading priority numbers from memory 608 such that the offset function is transparent to the user of DSP 602. For yet another embodiment, DSP 602 can detect that PNEW is equal to zero when an insert instruction is received by DSP 602, or when an underflow condition occurs in decrement circuit 3304. PNEW can then be incremented by one and be provided to decrement circuit 3304, or PNEW can be provided directly to priority index table 606 by multiplexer 3302. After all the priority numbers greater than zero are detected and updated by priority logic 610, then memory 608 can be searched for any existing entry that already has priority number 0. If so, this number can be incremented (i.e., by using the COUNT and UP signals, or by overwriting priority number 0 with priority number 1), and then PNEW added to memory 608.

A policy statement and its priority number may also be deleted or removed using priority logic 610. For example, policy statement 1708 may be removed from array 604, and its priority number 1 may be removed from memory 608 as follows. In response to a delete instruction, the priority number

(3308) to be deleted is provided to policy index table 606 via multiplexer 3302 (INST logic low), and is compared with the entries in the priority memory. For this embodiment, priority memory 608 is a CAM, or has one or more compare circuits external to the memory to perform the comparison. For an alternative embodiment, the policy statement associated with priority number 3308 may be

compared with the existing entries in CAM array 604. Since location 608₃ matches priority number 3308, both the policy statement and the priority number are invalidated or deleted. As described above, this may be accomplished by setting a valid bit(s) to an appropriate state for that row in priority memory 608. If the valid bit(s) is located in CAM array 604, and is

inaccessible by priority memory 608, then priority logic 610 can then identify the location of the matching entry (as described above), and provide the indication to IAD₀-IAD_{N-1}. These signal lines can then be provided through encoder 612 back to CAM array 604 (e.g., through a decoder) to address the appropriate location and invalidate the policy statement. Alternatively, the signal lines

IAD₀- IAD_{N-1} can be provided directly to CAM array 604 to access the appropriate location for invalidation. The delete process may stop at this point and potentially leave gaps in the priority numbers. Alternatively, the process may update the remaining entries in the priority memory. When the process continues, the deleted priority number 3308 is provided to policy index table 606

via multiplexer 3302, to determine which entries in the priority memory have a lower priority than the deleted priority number. Priority logic 610 identifies priority numbers 2 and 3 as having lower priorities (i.e., larger numerical

values), and these numbers are decremented. The resultant table entries are as shown in Figure 33.

The comparison and updating functions described with respect to Figures 33 and 34 may also be performed by priority logic 610 on priority numbers assigned in descending priority order. For this embodiment, decrement circuit 3304 may be replaced with an increment circuit that increments PNEW prior to supplying it to priority index table 606 for comparison as part of an insert function.

Note that if the largest priority number 2^n-1 (where n is the number of bits in a row of memory 608) is to be inserted, then the increment circuit would generate a number larger than can be handled by priority index table 606. For one embodiment, only priority numbers less than 2^n-1 may be used. For another embodiment, offset circuitry may be included in DSP 602 to subtract an offset value to PNEW before it is provided to the increment circuit and/or multiplexer 3302. The offset circuitry may also be used to add the offset value when reading priority numbers from memory 608 such that the offset function is transparent to the user of DSP 602. For yet another embodiment, DSP 602 can detect that PNEW is equal to 2^n-1 when an insert instruction is received by DSP 602, or when an overflow condition occurs in the increment circuit. PNEW can then be decremented by one and be provided to the increment circuit, or PNEW can be provided directly to priority index table 606 by multiplexer 3302. After all the priority numbers greater than 2^n-1 are detected and updated by priority logic 610, then memory 608 can be searched for any existing entry that already has priority number 2^n-1 . If so, this number can be decremented (i.e., by using the COUNT and DOWN signals, or by overwriting priority number 2^n-1 with priority number 2^n-2), and then PNEW added to memory 608.

The embodiments described above in Figures 33 and 34 include the ability of priority logic 610 to increment and/or decrement the priority numbers stored in priority memory 608. Priority memory 608 may be formed as counters that store the priority numbers. Figure 35 shows one embodiment of a row of priority memory 608 in which the priority logic 610 includes a separate priority logic element for each bit of a priority number. For this embodiment, the priority logic includes row 1201₀ of priority logic elements from Figure 12A, coupled to receive priority number bits stored in counter 2306₀ from Figure 23. PSUB bits PSUB_{n-1}-PSUB₀ are provided on signal lines 1208_{n-1}-1208₀, respectively, for comparison by compare circuits 1206_{n-1,0}-1206_{0,0} with the priority bits stored in memory elements 2314_{n-1,0}-2314_{0,0}. When BCT is asserted to a high logic state, logic 3306₀ drives match line segment 1210_{n-1,0} to a high logic state, and PSUB is provided on signal lines 1208. BCT is also provided to each of compare circuits 1206 such that the circuits do not provide their compare results back to signal lines 1208. Thus, only PSUB will be supplied on signal lines 1208. For an alternative embodiment, a separate set of signal lines may be used for the updating functions and the classification function. When BCT is asserted to a low logic state, ML₀ is provided to match lines segment 1210_{n-1,0} and the classification function of Figure 7 may be performed.

For this embodiment, the logical state of IAD₀ indicates the comparison result between PSUB and the priority number stored in counter 2306₀. If IAD₀ indicates that PSUB has a more significant priority than the priority number stored in counter 2306₀, then counter 2306₀ is incremented. If, however, IAD₀ indicates that PSUB is equal to or has a less significant priority number than the priority number in counter 2306₀, then counter 2306₀ is not incremented.

The operation of the embodiment of Figure 35 to update priority numbers in response to an insert instruction can be illustrated with the example shown in

Figure 36. In this example, a 3x3 matrix of rows and columns stores three priority numbers. For other embodiments, any numbers of rows and columns can be used. Row 0 stores priority number 010 having the decimal equivalent of the number 2, row 1 stores priority number 001 having the decimal equivalent of the number 1, and row 2 stores priority number 000 having the decimal equivalent of the number 0. For this example, the priority numbers are assigned in ascending priority order such that 000 is the most significant priority number, 001 is the next most significant priority number, and 010 is the least significant priority number.

When a new policy statement having new a priority number PNEW 001 is to be loaded into CAM array 604, BCT is asserted to a high logic state, the control signals on match line segments $1210_{2,2}$, $1210_{2,1}$, and $1210_{2,0}$ are asserted to a high logic state, and priority logic elements 1201_2 - 1201_0 compare PSUB 000 (PNEW decremented by 1) with their corresponding priority numbers stored in counters 2306_2 - 2306_0 , respectively. Priority logic 1201_2 determines that 000 is equal to PSUB; priority logic 1201_1 determines that 001 is greater than PSUB; and priority logic 1201_0 determines that 010 is greater than PSUB. Priority logic 1201_2 leaves unchanged the priority number stored in counter 2306_2 . Circuits 1201_1 and 1201_0 cause counters 2306_1 and 2306_0 to increment by one their existing priority numbers. The new policy statement and PNEW can then be loaded into CAM array 604 and priority memory 608, respectively.

The comparisons of the most significant bit $PSUB_2$ with the most significant bits stored in counters 2306_2 - 2306_0 are resolved first. When an individual compare circuit determines that its corresponding memory storage element stores the same logic state as the PSUB bit, the compare circuit leaves the control signal of the next match line segment unaffected such that it has the same logical state of the previous match line segment. Since each of memory storage

elements 2314_{2,2}, 2314_{2,1}, and 2314_{2,0} store a logic zero and PSUB₂ is a logic zero, the control signals on match line segments 1210_{1,2}, 1210_{1,1}, and 1210_{1,0} remain asserted to a high logic state.

The comparison of the next most significant bit PSUB₁ with priority number bits stored in memory elements 2314_{1,0}, 2314_{1,1} and 2314_{1,2} is then resolved. Since memory elements 2314_{1,2} and 2314_{1,1} store the same logic states as PSUB₁, the control signals on match line segments 1210_{0,2} and 1210_{0,1} remain in a high logic state. Compare circuit 1206_{1,0}, however, determines that the priority number stored in counter 2306₀ is greater than PSUB because: (1) memory element 2314_{1,0} stores a logic one; (2) match line segment 1210_{1,0} is a high logic state; and (3) PSUB₁ is a logic zero on signal line 1208₁. When compare circuit 1206_{1,0} makes this determination, it discharges match line segment 1210_{0,0} to a low logic state. When a match line segment is discharged, all subsequent match line segments in that row are discharged to a low logic state such that IAD for that row is discharged to a low logic state. For this embodiment, when IAD is a low logic state, the counter associated with that row will be incremented when UP is asserted. Thus, IAD₀ will cause counter 2306₀ to be incremented when UP is asserted.

Lastly, the comparison of the least significant bit PSUB₀ with priority number bits stored in memory elements 2314_{0,0}, 2314_{0,1} and 2314_{0,2} is resolved. As indicated above, priority logic 1201₀ has already determined that 010 is greater than 000 and thus IAD₀ remains asserted to a low logic state. Since the control signal on match line segment 1210_{0,2} is in a high logic state, compare circuit 1206_{0,2} compares the logic zero stored in memory element 2314_{0,2} with the logic zero of PSUB₀, and determines that PSUB is the same number as that stored in counter 2306₂. In response, IAD₂ remains in a high logic state such that counter 2306₂ is not updated. Compare circuit 1206_{0,1}, however, determines that PSUB₀ is

less than the logic one stored in memory element 2314_{0,1}. Since the control signal is high on match line segment 1210_{0,1}, compare circuit 1206_{0,1} discharges IAD1 such that the priority number 001 stored in counter 2306₁ will be incremented. After a sufficient amount of time has passed such that all of the bit comparisons
5 have had time to resolve, the UP signal can be asserted such that the priority numbers in counters 2306₁ and 2306₀ are incremented. For one embodiment, the IAD signals can be latched prior to providing the UP signal. Subsequently, PNEW can be added to memory 608.

Any compare circuits may be used for compare circuits 1206 to implement
10 the process illustrated above. For example, the compare circuits illustrated above that implement Table 1 can be used with additional control for the BCT control signal. One embodiment of the modified circuit of Figure 14 is shown in Figure 37. In Figure 37, transistor 1416 is included in series with transistors 1406 and 1408 between signal line 1208_{n-1} and ground. The gate of transistor 1416 is
15 controlled by BCT via inverter 1420. A similar modification can be made to any of the other embodiments of the compare circuits 2314. OR gate 1418 is one embodiment of logic 3302.

While Figure 36 shows an embodiment of inserting a new priority number into memory 608, the priority numbers can be decremented for a delete
20 function by asserting the DOWN signal after all of the IAD signals are resolved.

The embodiment of Figure 35 can also insert and delete priority numbers assigned in descending priority order. Figure 38 shows such an example with of a 3x3 matrix in which row 0 stores priority number 111 having the decimal equivalent of the number 7, row 1 stores priority number 110 having the decimal
25 equivalent of the number 6, and row 2 stores priority number 101 having the decimal equivalent of the number 5. Thus, 111 is the most significant priority

number, 110 is the next most significant priority number, and 101 is the least significant priority number.

When a new policy statement having a new priority number PNEW 110 is to be loaded into CAM array 604, BCT is asserted to a high logic state, control signals on match line segments 1210_{2,2}, 1210_{2,1}, and 1210_{2,0} are asserted to a high logic state, and priority logic elements 1201₂-1201₀ compare PSUB 111 (PNEW incremented by 1) with their corresponding priority numbers stored in counters 2306₂-2306₀, respectively. Priority logic 1201₂ determines that 111 is equal to PSUB; priority logic 1201₁ determines that 110 is less than PSUB; and priority logic 1201₀ determines that 101 is less than PSUB. Priority logic 1201₂ leaves unchanged the priority number stored in counter 2306₂. Circuits 1201₁ and 1201₀ cause counters 2306₁ and 2306₀ to decrement by one their existing priority numbers. The new policy statement and PNEW can then be loaded into CAM array 604 and priority memory 608, respectively.

The comparisons of the most significant bit PSUB₂ with the most significant bits stored in counters 2306₂-2306₀ are resolved first. When an individual compare circuit determines that its corresponding memory storage element stores the same logic state as the PSUB bit, the compare circuit leaves the control signal of the next match line segment unaffected such that it has the same logical state of the previous match line segment. Since each of memory storage elements 2314_{2,2}, 2314_{2,1}, and 2314_{2,0} store a logic one and PSUB₂ is a logic one, the control signals on match line segments 1210_{1,2}, 1210_{1,1}, and 1210_{1,0} remain asserted to a high logic state.

The comparison of the next most significant bit PSUB₁ with priority number bits stored in memory elements 2314_{1,0}, 2314_{1,1} and 2314_{1,2} is then resolved. Since memory elements 2314_{1,2} and 2314_{1,1} store the same logic states as PSUB₁, the control signals on match line segments 1210_{0,2} and 1210_{0,1} remain in a

high logic state. Compare circuit 1206_{1,0}, however, determines that the priority number stored in counter 2306₀ is less than PSUB because: (1) memory element 2314_{1,0} stores a logic zero; (2) match line segment 1210_{1,0} is a high logic state; and (3) PSUB₁ is a logic one on signal line 1208₁. When compare circuit 1206_{1,0} makes this determination, it discharges match line segment 1210_{0,0} to a low logic state. When a match line segment is discharged to a low logic state, all subsequent match line segments in that row are discharged such that IAD for that row is discharged to a low logic state. When IAD is a low logic state, the counter associated with that row will be decremented when DOWN is asserted. Thus, IAD₀ will cause counter 2306₀ to be decremented when DOWN is asserted.

Lastly, the comparison of the least significant bit PSUB₀ with priority number bits stored in memory elements 2314_{0,0}, 2314_{0,1} and 2314_{0,2} is resolved. As indicated above, priority logic 1201₀ has already determined that 101 is less than 111 and thus IAD₀ remains discharged to a low logic state. Since the control signal on match line segment 1210_{0,2} is in a high logic state, compare circuit 1206_{0,1} compares the logic one stored in memory element 2314_{0,2} with the logic one of PSUB₀, and determines that PSUB is the same number as that stored in counter 2306₂. In response, IAD₂ remains in a high logic state such that counter 2306₂ is not updated. Compare circuit 1206_{0,1}, however, determines that PSUB₀ is greater than the logic zero stored in memory element 2314_{0,1}. Since the control signal is asserted on match line segment 1210_{0,1}, compare circuit 1206_{0,1} discharges IAD₁ such that the priority number 110 stored in counter 2306₁ will be decremented. After a sufficient amount of time has passed such that all of the bit comparisons have had time to resolve, the DOWN signal can be asserted such that the priority numbers in counters 2306₁ and 2306₀ are decremented. For one embodiment, the IAD signals can be latched prior to providing the DOWN signal. Subsequently, PNEW can be added to memory 608.

Any compare circuits may be used for compare circuits 1206 to implement the process illustrated in Figure 38. For example, the compare circuits illustrated above that implement Table 3 can be used with additional control for the BCT control signal as illustrated in Figure 37. While Figure 38 shows an embodiment of inserting a new priority number into memory 608, the priority numbers can be incremented for a delete function by asserting the UP signal after all of the IAD signals are resolved.

DSP 602 can perform the updating functions described above with respect to Figures 17-38 asynchronously or synchronously. When DSP 602 operates synchronously, it receives a clock signal that may be used to clock in the policy search key and an instruction that causes the updating functions to be performed by DSP 602. DSP 602 may implement the updating functions in one or more clock cycles.

Depth Cascading DSP Devices

As described above, DSP 602 stores policy statements in CAM array 604 and identifies the highest priority matching policy statement without having to presort or prearrange the policy statements in the CAM array. DSP 602 may also be included in a system that has multiple DSPs connected in a depth cascade configuration that expands the number of memory locations in CAM array 604 and priority memory 608 to accommodate more policy statements and their associated priority numbers.

Figure 39 shows a CAM system 3900 that includes three DSPs 3902, 3904, and 3906 interconnected in a depth cascade configuration. Policy statements may be arbitrarily stored in DSPs 3902-3906 without initially presorting or prearranging the statements. When a classification or filtering operation, such as that described in Figure 7, is performed by each of the DSPs in system 3900, the DSP that stores the highest priority matching statement for system 3900, as a

whole, may be determined as will be described in more detail below. The DSP that includes the highest priority matching statement for system 3900 may be determined regardless of the number of CAM locations filled or occupied in each CAM array of each DSP (i.e., regardless of whether one or more CAM
5 arrays of DSPs 3902-3906 is full, empty, or partially full).

Any number of DSPs may be depth cascaded as shown in Figure 39. The total memory size of system 3900 is larger than the individual memory sizes of each of the CAM arrays in each DSP 3902-3906. For example, if each CAM array is a 16k x 144 bit CAM array, then system 3900 may operate as a DSP that
10 includes a 48k x 144 bit CAM array. Any other size of CAM arrays may be used. Also, CAM arrays of different widths may be depth cascaded together. Additionally, system 3900 may be formed from stand-alone DSPs, or may be formed from integrated circuits on a common substrate.

Each DSP receives in parallel a clock signal CLK, policy statements and
15 priority numbers on data bus DBUS, and instructions on instruction bus IBUS. For alternative embodiments, the policy statements and priority numbers may be provided on different buses. Other input signals may also be simultaneously provided to each of DSPs 3902-3906 including word enable signals, reset signals, chip enable signals, and the like. DSPs 3902-3906 also output their external
20 address signal to a common HPM bus. When performing the classification function of Figure 7, the DSP that has the highest priority matching policy statement will take control of the HPM bus. Each DSP 3902-3906 may include any of the embodiments of DSP 602 described above.

DSPs 3902-3906 may each include a cascade down input /CDI, a cascade
25 down output /CDO, a cascade up input /CUI, a cascade up output /CUO, priority number down inputs PNDI, priority number down outputs PNDO, priority number up inputs PNUI, and priority number up outputs PNUO. The

PNDO and PNUO outputs provide the most significant priority number PNUM generated by the priority logic in each priority index table in each of DSP 3902-3906. The cascade up and down outputs provide signals that indicate when PNUM is valid on the PNDO and PNUO outputs. For other embodiments, the cascade inputs and outputs may be omitted and CLK or match flag up and down signals may be used to validate the PNUM outputs. For an alternative embodiment, match flag up and down signals generated by flag logic coupled to CAM array 604 may be used to validate the PNUM outputs. Additional outputs or inputs such as full flag up and down pins may also be included.

Each CAM DSP generally has its /CDO output coupled to the /CDI input of the next DSP, its /CUO output coupled to the /CUI of the previous DSP, its /CDI input coupled to the /CDO of the previous DSP, its /CUI input coupled to the /CUO of the next DSP, its PNDO outputs coupled to the PNDI inputs of the next DSP, its PNDI inputs coupled to the PNDO outputs of the previous DSP, its PNUI inputs coupled to the PNUO outputs of the next DSP, and its PNUO outputs coupled to the PNUI inputs of the previous DSP.

DSP 3902 may be designated as the DSP that has the lowest logical system addresses by connecting its /CDI input to ground, and its PNDI inputs to VDD. For another embodiment, the PNDI inputs may be connected to ground and/or the /CDI inputs of DSP 3902 may be connected to VDD. DSP 3902 has the lowest logical or numerical addresses of, for example, zero to N-1. DSP 3904 has logical addresses N to M-1, and DSP 3906 has logical addresses M to W-1, where system 3900 has a total of W ternary CAM words available for storing policy statements. DSP 3906 may be designated as the last DSP in the cascade by connecting its PNUI inputs to VDD, and its /CUI input to ground. For another embodiment, the PNUI inputs of DSP 3906 may be connected to ground, and/or the /CUI input may be connect to VDD.

09316778-033401
The operation of system 3900 for implementing the classification or filtering function of Figure 7 is as follows. Initially, one or more of DSPs 3902-3906 is loaded with policy statements and corresponding priority numbers. In response to an instruction provided on the IBUS and the policy search key provided on the DBUS, each DSP 3902-3906 compares the policy search key with the policy statements stored in its CAM array 604 (step 702). The priority logic in each priority index table identifies its most significant priority number PNUM associated with one of its matching policy statements. Each DSP also determines the address of its PNUM in its priority memory 608 (steps 706 and 708). Each DSP then compares the priority numbers provided on its PNDI and PNUI pins with its own PNUM to determine whether it has a more significant priority number. If a DSP has a more significant or equal priority number than that provided on its PNUI pins, then the DSP outputs its own PNUM to its PNUO pins. If, however, a DSP has a lower priority PNUM than those provided on the PNUI pins, then the DSP outputs the priority number from its PNUI pins to its PNUO pins. Similarly, if a DSP has a more significant priority number than that provided on its PNDI pins, then the DSP outputs its own PNUM to its PNDO pins. If, however, a DSP has an equal or lower priority PNUM than those provided on the PNDI pins, then the DSP outputs the priority number from its PNDI pins to its PNDO pins.

By simultaneously providing and comparing priority numbers both down and up through system 3900, each DSP will ultimately determine whether it has the most significant priority number in the entire system. When a device has the most significant priority number for the system, it may take control of the HPM bus and output address information to access a location in route memory 614. For this embodiment, route memory 614 is as deep as the number of CAM memory location in system 3900.

Each of DSPs 3902-3906 can also update their own policy statement tables and priority memories as needed in response to an insertion or deletion instruction provided to system 3900. A new policy statement and priority number, for example, may be loaded into the DSP that has the next free address of system 3900 as determined by full flag signals (not shown).

Figure 40 shows DSP 4000 that is one embodiment of DSPs 3902-3906. DSP 4000 is any of the embodiments of DSP 602 that further includes cascade logic 4004, output buffer 4002, instruction decoder 4006, read and write (R/W) circuits 4008 and 4010, and flag logic 4016. DSP 4000 may also include registers to store the policy statements and priority numbers prior to use by CAM array 604 and priority index table 606. Cascade logic 4004 is coupled to the /CDI, /CUI, PNDI, and PNUI inputs, and to the /CDO, /CUO, PNDO, and PNUO outputs. For another embodiment, the cascade inputs and outputs (i.e., /CDI, /CDO, /CUI, and /CUO) may be omitted and thus not coupled to cascade logic 4004. Cascade logic 4004 may also receive and output match flag and full flag cascade signals.

Cascade logic 4004 receives a search signal SCH on line 4012 from instruction decoder 4006 indicating that the classification or filtering operation will be performed by DSP 4000. Cascade logic 4004 may also receive a match flag signal /MF from flag logic 4016 indicating whether CAM array 604 has an address that matches a policy search key. Cascade logic 4004 compares the priority number on its PNDI inputs with its own most significant priority number PNUM from priority index table 606, and outputs the more significant number from among these two priority numbers to its PNDO outputs. Similarly, cascade logic 4004 compares the priority number on its PNUI inputs with its own PNUM, and outputs the more significant number from among these two priority numbers to its PNUO outputs. If DSP 4000 has the most significant

priority number for the entire system 3900, then it asserts its output buffer control signal OBCNTL on signal line 4014 to enable output buffer 4002 to take control of the HPM bus.

Figure 41 is a block diagram of cascade logic 4100 that is one embodiment of cascade logic 4004 of Figure 40. Cascade logic 4100 includes cascade down logic 4104 and cascade up logic 4102. Cascade down logic 4104 generates signals for the /CDO and PNDO outputs in response to /CDI, PNDI inputs, and SCH. Cascade up logic 4102 generates signals on the /CUO and PNUO outputs in response to /CUI, PNUI inputs, and SCH. Other embodiments of cascade logic 4004 may be used. For example, match flag signals may be used.

Cascade down logic 4104 includes signal generator 4106 and compare circuit 4110. Signal generator 4106 receives CLK and SCH from instruction decoder 4006. SCH indicates that the classification of filtering function is to be performed by DSP 4000. Signal generator 4106 asserts /CDO to a low logic state after /CDI is asserted to a low logic state and when the priority number on the PNDO outputs is valid. Signal generator 4106 may be any logic circuit that performs this function. For an alternative embodiment, SCH may be omitted and signal generator 4106 may generate /CDO in response to CLK only. Compare circuit 4110 compares the internal most significant priority number PNUM with the priority number from the PNDI inputs. If the priority number on the PNDI inputs is equal to or more significant than PNUM, then compare circuit 4110 outputs the priority number from its PNDI inputs to its PNDO outputs and asserts signal line 4116 to a low logic state. If, however, PNUM is more significant than the priority number on the PNDI inputs, then compare circuit 4110 outputs PNUM to its PNDO outputs and asserts line 4116 to a high logic state.

Cascade up logic 4102 includes signal generator 4108 and compare circuit 4112. Signal generator 4108 receives CLK and SCH from instruction decoder 4006. Signal generator 4108 asserts /CUO to a low logic state after /CUI is asserted to a low logic state and when the priority number on the PNUO outputs is valid. Signal generator 4108 may be any logic circuit that performs this function. For an alternative embodiment, SCH may be omitted and signal generator 4108 may generate /CUO in response to CLK only. Compare circuit 4112 compares PNUM with the priority number from the PNUI inputs. If the priority number on its PNUI inputs is more significant than PNUM, then compare circuit 4112 outputs the priority number from its PNUI inputs to its PNUO outputs and asserts signal line 4118 to a low logic state. If, however, PNUM is equal to or more significant than the priority number on its PNUI inputs, then compare circuit 4112 outputs PNUM to the PNUO outputs and asserts line 4118 to a high logic state. When signal lines 4116 and 4118 are both at high logic states, AND gate 4114 asserts OBCNTL on line 4014 to a high state to enable output buffer 4002 to take control of the HPM bus.

For an alternative embodiment in which the cascade inputs (/CDI and /CUI) and cascade outputs (/CDO and /CUO) are omitted, signal generators 4106 and 4108 may also be omitted. For yet another embodiment in which multiple /CDI-/CDO and /CUI-/CUO pins are included, multiple pairs of signal generators may also be included, and/or match flag signals may be included.

Classless Inter Domain Routing (CIDR)

DSP 602 can also be used to process Internet Protocol (IP) packets that use the Classless Inter Domain Routing (CIDR) scheme. With CIDR, an IP address has a generalized network prefix of a particular number bits of 32-bit IPv4 or a 128-bit IPv6 address. The network prefix or mask indicates the number of left-

most contiguous bits in the IP address that are used to filter an IP address in a routing table. That is, the network prefix indicates the number of higher-order or left-most contiguous bits in the IP address that participate in an address comparison with the routing table.

5 Conventional ternary CAM devices such as CAM 300 of Figure 3 can store the IP addresses in rows 305, and their corresponding prefixes in rows 310.

Routing information associated with a particular IP address is loaded into a corresponding address location in route memory 308. Due to the operation of priority encoder 306, IP addresses are generally presorted or prearranged prior
10 to entry into a CAM device such that the IP address with the longest network prefix is located in the lowest logical address of the CAM array, and the IP address with the shortest network prefix is located in the highest logical address of the CAM array. When the IP addresses are presorted, a search on the CAM array for a particular IP address will identify the IP address that has the longest
15 corresponding prefix, that is, will identify the best match.

A considerable amount of time is generally required to prearrange all of the CIDR address entries prior to loading the entries into a CAM device.

Additionally, a considerable amount of time and overhead is also generally required to maintain the order of the routing table when entries are deleted or
20 overwritten, or when new entries are to be added. Other architectures have been proposed that increase the size of the CAM array by adding additional logic in the CAM array itself and another match coupled to the rows of mask cells.

DSP 602 of Figure 6 can be used to process IP addresses without adding additional logic or signal lines to ternary CAM array 604. IP addresses can be
25 loaded into CAM cell rows 605, and the corresponding decoded prefix data can be loaded into mask rows 607. Decoding logic can be provided in DSP 602 to decode the prefix number. Additionally, the prefix data is encoded into a binary

number and stored in corresponding locations in priority memory 608. Encoding logic can be provided in DSP 602 to encode the prefix number into a binary (or other code) number.

When a search is performed for the IP address with the longest prefix, all matching locations in CAM array 604 will assert their corresponding match lines ML_0 - ML_{N-1} . Priority logic 610 then compares, with each other, the encoded prefix numbers associated with the matching IP address. Priority logic 610 identifies the most significant encoded prefix number (i.e., the highest prefix number), and identifies its location in priority memory 608 to IAD_0 - IAD_{N-1} . The encoded most significant prefix number may also be output from DSP 600. Encoder 612 then encodes the identified location into an address for output to HPM bus 616. The corresponding route information can then be accessed in route memory 614. As in the previous embodiments described above, route memory 614 may also be included within DSP 602.

For another embodiment, there may more than one identical most significant priority numbers identified by priority logic 610. For this embodiment, encoder 612 may be a conventional priority encoder that determines which address to output based on a predetermined priority (i.e., based on logical address locations).

DSP 602 can process the CIDR based IP addresses without preloading the IP addresses in the CAM array in a predetermined order. Additionally, new IP address may be added at the next free address or any other designated address in CAM array 604 without reordering or reloading the CAM array. This can reduce the hardware and/or software needed for table management of the CAM array, and can increase the performance of a router incorporating the DSP.

The operation of DSP 602 for processing CIDR based IP addresses can be illustrated by the example of Figure 42. In Figure 42, ternary CAM array 604 has

IP address 168.0.0.0/8 stored at location 604₀, 168.69.0.0/16 stored at location 604₁, and 168.69.62.0/24 stored at location 604₂. For this embodiment, each IP address is stored in array 604 as four eight-bit binary numbers. Also for this embodiment, when the decoded prefix data is a logic zero it does not mask the
5 corresponding bits of the IP address. Priority memory 608 stores the prefixes 8, 16, and 24 at locations 608₀, 608₁, and 608₂, and the corresponding routing information RI₀, RI₁, and RI₂ are stored at locations 614₀, 614₁, and 614₂ of route memory 614.

A search key of 168.69.43.100 is provided to CAM array 604, and the IP
10 address with the best match (i.e., the longest prefix data) is determined as follows. When the search key is compare with the IP addresses, 168.69.0.0/16 and 168.0.0.0/8 are both identified as matches and ML₀ and ML₁ asserted. Between these two IP addresses, 168.69.0.0/16 is the best match as it has a longer prefix. Priority logic 610 compares the prefixes 16 and 8 stored at locations 608₀
15 and 608₁ and determines that 16 is greater than 8. The priority logic outputs 16 as the longest matching prefix to PNUM, and also identifies location 608₁ by asserting IAD₁. Encoder 612 then encodes IAD₀-IAD_{N-1} and generates an address of 1 on HPM bus 616 to access route information RI₁ in route memory 614.

Any of the embodiments of DSP 602 described above can be used to
20 implement priority logic 610 to process CIDR based IP addresses and their prefix numbers. Additionally, multiple DSPs can be depth cascaded as described in Figures 39-41 to implement a system that provides a sufficient number of CAM array locations for storing IP addresses.

Priority Encoding In A Segmented Classification System

When the policy statement table 604 and the priority index table 606 of the classification system 600 (all shown in Figure 6) are implemented using large memory arrays, resistance-capacitance (RC) delay on the arrays' signal lines may
5 cause undesirable loading affects. According to an embodiment of the present invention, the classification system 600 shown in Figure 6 may be segmented to reduce the loading affects associated with RC delay.

Figure 43 illustrates a first embodiment of a segmented classification system 4300 according to the present invention. The segmented classification
10 system 4300 includes a plurality of data classification blocks where block 4310 represents a first data classification block, block 4320 represents a second data classification block, and block 4330 represents an Nth data classification block where N can be any number. Each of the data classification blocks 4310, 4320, and 4330 includes a segmented portion of the policy statement table 604, priority
15 index table 606, and encoder 612 (shown in Figure 6) of the classification system 600 or the systems shown in Figures 17, 33, 39 and 42. It should be appreciated that the classification system 600 may be segmented into any number of data classification blocks.

Each of the data classification blocks 4310, 4320, and 4330 includes a data
20 table 4311, 4321, and 4331, respectively. Each of the data tables 4311, 4321, and 4331 is a segmented portion of the policy statement table 604. Similar to the policy statement table 604, each of the data tables 4311, 4321, and 4331 may be implemented using a ternary CAM array, binary CAM array or any other type of memory to store data. The data that is stored in the data tables 4311, 4321, and
25 4331 may include, for example, policy statements that are compared with processed policy information of an incoming packet. In an example where policy statement table 604 includes three data classification blocks 4310, 4320, and 4330 that stores policy statements 0 to z, data table 4311 may store policy

statements 0 to x, data table 4321 may store policy statements x+1 to y, and data table 4331 may store policy statements y+1 to z.

Each of the data classification blocks 4310, 4320, and 4330 includes a priority index table 4312, 4322, and 4332, respectively. Each of the priority index tables 4312, 4322, and 4332 may be coupled to a data table in its data classification block via a plurality of match lines (not shown). Each of the priority index tables 4312, 4322, and 4332 is a segmented portion of the priority index table 606. Similar to the priority index table 606, each of the priority index tables, 4312, 4322, and 4332 includes a priority memory (not shown) that stores priority numbers assigned to policy statements in the data table in its data classification block. The priority memory in priority index table 4312 may store priority numbers for policy statements 0 to x. The priority memory in priority index table 4322 may store priority numbers for policy statements x+1 to y. When system 4300 includes only three data classification blocks, the priority memory in priority index table 4332 may store priority numbers for policy statements y+1 to z. Each of the priority index tables 4312, 4322, and 4332 also includes priority logic (not shown). Each priority logic compares the priority numbers for its corresponding policy statements stored in its data classification block that match an incoming packet to determine a most significant priority number in that block. This most significant priority number is given the designation of being the block priority number (BPN) for its data classification block. Each priority index table 4312, 4322, and 4332 may also include inequality circuits as shown in Figure 17 and/or decrement and multiplexer circuitry as shown in Figure 33.

The segmented classification system 4300 includes a device index processor 4360 that determines a device index (DIN) that indicates the index or address of an entry in one of the data tables 4311, 4321, or 4331 that matches a

search key and has associated with it the most significant priority number. The device index processor 4360 receives a block priority number for each of the data classification blocks 4310, 4320, and 4330. The device index processor 4360 determines a most significant block priority number from among the block
5 priority numbers received. The device index processor 4360 outputs an enable signal (EN) to each of the data classification blocks 4310, 4320, and 4330 that indicates whether or not a data classification block originates the most significant block priority number.

Each of the data classification blocks 4310, 4320, and 4330 includes a row
10 enable logic circuit 4313, 4323, and 4333, respectively. Each row enable logic circuit 4313, 4323, and 4333 may be coupled to a priority index table in its data classification block via a plurality of internal address lines from the priority index table. Each of the plurality of internal address lines from a priority index table corresponds to a priority number stored in the priority index table. When
15 a priority number is a block priority number, its corresponding internal address line carries a match signal from the priority index table that indicates a match or mismatch state. Each row enable logic circuit 4313, 4323, and 4333 may be configured to propagate a match signal from the priority index table in its data classification block in response to receiving an enable signal from the device
20 index processor 4360 that indicates that its data classification block originates the most significant block priority number. When, however, an enable signal from the device index processor 4360 indicates that a corresponding data classification block does not originate the most significant block priority number, the corresponding row enable logic circuit causes the match signals on the internal
25 address lines to indicate mismatch states.

According to an embodiment of the data classification system 4300, each row enable logic circuits 4313, 4323, and 4333 is implemented using a plurality of

circuits that perform an AND function. Each of the circuits may have a first input coupled to one of the internal address lines from the priority index table in its data classification block and a second input from an enable signal from the device index processor 4360. The output of the circuits may be transmitted onto an internal address line to an encoder in its data classification block. It should be appreciated that the row enable logic circuits 4313, 4323, and 4333 may be implementing using any known circuitry or technique.

Each of the data classification blocks 4310, 4320, and 4330 includes a match flag generator 4314, 4324, and 4334, respectively. Each match flag generator 4314, 4324, and 4334 may be coupled to the row enable logic circuit in its data classification block via the plurality of internal address lines from the row enable logic circuit. The match flag generator detects whether one of the internal address lines from the row enable logic circuit in its data classification block indicates a match. Each match flag generator 4314, 4324, and 4334 may be configured to generate a match flag signal (MF) that indicates whether one of the internal address lines from the row enable logic circuit in its data classification block indicates a match. An internal address line from the row enable logic circuit is asserted when its data classification block originates a block priority number that is the most significant block priority number. According to an embodiment of the present invention, the match flag signal may be an asserted or de-asserted signal line.

According to an embodiment of the data classification system 4300, each match flag generator 4314, 4324, and 4334 is implemented using a circuit that performs an OR function. Each circuit may have inputs coupled to the internal address lines from the row enable logic circuit from its data classification block. The output of each of the circuits is transmitted to the device index processor as

a match flag signal. It should be appreciated that the match flag generators 4314, 4324, and 4334 may be implemented using any known circuitry or technique.

Each of the data classification blocks 4310, 4320, and 4330 includes an encoder 4315, 4325, and 4335, respectively. Each encoder 4315, 4325, and 4335 may be coupled to a row enable logic circuit in its data classification block via a plurality of internal address lines from the row enable circuit. The identified location of the block priority number in the priority memory of each data classification block 4310, 4320, and 4330 is provided to its corresponding encoder 4315, 4325, and 4335, respectively via the internal address lines. The address or index of the most significant priority matching policy statement, relative to the data classification block, is determined by the encoder. This address or index is given the designation of being the block index (BIN) for its data classification block. When there is no match in a data table, the corresponding encoder may generate a default index.

The device index processor 4360 receives a match flag signal and block index for each of the data classification blocks 4310, 4320, and 4330. The device index processor 4360 determines a device address or device index (DIN) of the most significant priority matching policy statement stored in one of the data tables. The device index may be used to access corresponding route information from a memory such as route memory 614 shown in Figure 6.

Figure 43 illustrates the row enable logic circuits 4313, 4323, and 4333 positioned between priority index tables and encoders in their data classification blocks. It should be appreciated that in an alternative embodiment of the data classification system 4300, the row enable logic circuits 4313, 4323, and 4333 may be positioned to the right of the encoders in their data classification blocks.

Figure 44 illustrates a segmented data classification system 4400 that includes a device index processor 4460 that is one embodiment of the device

index processor 4360 shown in Figure 43. The device index processor 4460 includes a compare logic 4410, inhibit signal generator 4420, and address generator 4430. The compare logic 4410 may receive block priority numbers from each of the data classification blocks 4310, 4320, and 4330. From the block
5 priority numbers received from data classification blocks 4310, 4320, and 4330, the compare logic 4410 determines a most significant block priority number. According to an embodiment of the present invention, the compare logic 4410 may be configured to determine that the block priority number with either the highest or lowest numerical value is the most significant block priority number.

10 The compare logic 4410 may be implemented using, the bit-wise compare circuitry described in U.S. Application No. 09/439,834 filed on November 12, 1999 and entitled "METHOD AND APPARATUS FOR DETERMINING A LONGEST PREFIX MATCH IN SEGMENTED CONTENT ADDRESSABLE MEMORY DEVICE", herein incorporated by reference. For another
15 embodiment, compare logic 4410 may be implemented as priority index table 606 of Figure 6 including a priority memory to store the block priority numbers and priority logic to compare the block priority numbers. It should be appreciated that the compare logic 4410 may be implemented using any other known circuitry or techniques.

20 The device index processor 4460 includes an inhibit signal generator 4420. The inhibit signal generator 4420 receives the most significant block priority number (MSBPN) from the compare logic 4410 and block priority numbers from the data classification blocks 4310, 4320, and 4330. According to an embodiment of the present invention, the inhibit signal generator 4420 compares the block
25 priority numbers from the data classification blocks 4310, 4320, and 4330 with the most significant block priority number to determine which data classification block originates the most significant block priority number. The inhibit signal

generator 4420 generates an enable signal to each of the data classification blocks 4310, 4320, and 4330 that indicates whether the data classification block originates the most significant block priority number. According to an embodiment of the device index processor 4460, an enable signal may be an asserted or de-asserted signal line.

Figure 45 illustrates an inhibit signal generator 4520 that is one embodiment of the inhibit signal generator 4420 shown in Figure 44. The inhibit signal generator 4520 includes a plurality of comparators 4510, 4511, and 4512. Each comparator compares the most significant block priority number from the compare logic 4410 (shown in Figure 44) with a block priority number from one of the data classification blocks. Comparator 4510 compares the most significant block priority number with a block priority number from data classification block 4310 (shown in Figure 44). Comparator 4511 compares the most significant block priority number with a block priority number from data classification block 4320 (shown in Figure 44). Comparator 4512 compares the most significant block priority number with a block priority number from data classification block 4330 (shown in Figure 44). Each comparator 4510, 4511, and 4512 generates an enable signal that is transmitted to the data classification block that originated the compared block priority number. The enable signal indicates whether the data classification block originates the most significant block priority number. It should be appreciated that the comparators 4510-4512 may be implemented using any known circuitry or technique. In an embodiment of the device index processor 4460 where the compare logic 4410 is implemented using a priority index table, the inhibit signal generator may be omitted. For this embodiment, match line segments corresponding to a stored block priority number in the priority index table of compare logic 4410 may provide the enable

signals corresponding to a row enable logic circuit of a data classification block that originates the block priority number.

Referring back to Figure 44, the device index processor 4460 includes an address generator 4430. The address generator receives match flag signals and
5 block indices (BIN 0-BIN N) from each of the data classification blocks 4310, 4320, and 4330. The address generator 4430 processes the match flag signals and the block indices and generates a device index that corresponds to the most significant block priority number.

Figure 46 illustrates an address generator 4630 that is one embodiment of
10 the address generator 4430 shown in Figure 44. The address generator 4630 includes a priority encoder 4610. The priority encoder 4610 operates similarly to the priority encoder 306 described in Figure 3. The priority encoder 4610 receives a match flag signal from each of the data classification blocks 4310, 4320, and 4330 (shown in Figure 44) that indicate whether the data classification block
15 originates the most significant block priority number. The priority encoder 4610 determines and outputs a block identifier corresponding to a data classification block that originates the most significant block priority number. When a plurality of data classification blocks originate the most significant block priority number, the priority encoder 4610 outputs a block identifier corresponding to a
20 data classification block having the highest priority among the plurality of data classification blocks. According to an embodiment of the present invention, data classification block 4310 has the highest priority, data classification block 4320 has the second highest priority, and data classification block 4330 has the Nth highest priority. The block identifier that is output from the priority encoder
25 4610 is given the designation of most significant block identifier (MSBID).

The address generator 4630 includes an index storage unit 4640. The index storage unit 4640 includes a plurality of storage elements (not shown).

Each of the plurality of storage elements has an address. When an address of a storage element is asserted, the index storage unit 4640 outputs the content of the storage element. The index storage unit 4640 receives block indices from the data classification blocks 4310, 4320, and 4330. According to an embodiment of the address generator 4630, the index storage unit 4640 stores the block indices received from the data classification blocks 4310, 4320, and 4330 in assigned storage elements for the data classification blocks 4310, 4320, and 4330.

The address generator 4630 includes a decoder 4620. The decoder 4620 receives the most significant block identifier from the priority encoder 4610. The decoder 4620 decodes the most significant block identifier into an address to access data in the index storage unit 4640. According to an embodiment of the address generator 4630, the decoder 4620 generates an address of a storage element in the index storage unit 4640 that stores a block index from a data classification block that corresponds to the most significant block identifier. When the decoded address is asserted, the index storage unit 4640 outputs a block index that is given the designation of the most significant block index (MSBIN).

For another embodiment, the index storage unit 4640 may be replaced by a multiplexer and decoder 4620 may generate select signals that cause the multiplexer to output a block index from a data classification block that corresponds to the most significant block identifier.

The address generator 4630 includes an index generator 4650. The index generator 4650 receives the most significant block identifier from the priority encoder 4610 and the most significant block index from the index storage unit 4640. According to an embodiment of the address generator 4630, the index generator 4650 generates a device index by designating the most significant block identifier as the most significant bits of the device index and designating

the most significant block index as the least significant bits of the device index.
Alternatively, index generator 4650 may be omitted.

According to an embodiment of the segmented data classification system 4400, the data classification blocks 4310, 4320, and 4330, and the device index
5 processor 4460 all reside on a single semiconductor substrate.

Fig. 47 flow chart illustrating a first method for generating a device index
in a segmented data classification system according to an embodiment of the
present invention. At step 4701, a first block priority number is determined.
The first block priority number is determined from a first plurality of priority
10 numbers. Each of the first plurality of priority numbers is associated with data
stored in a first data classification block (DCB) that matches a search key.

At step 4702, a second block priority number is determined. The second
block priority number is determined from a second plurality of priority
numbers. Each of the second plurality of priority numbers is associated with
15 data stored in a second data classification block that matches a search key. Steps
4701 and 4702 may be performed sequentially or simultaneously.

At step 4703, a most significant block priority number is determined. The
most significant block priority number is a block priority number having the
most significant value. The most significant value may be, for example, the
20 lowest or highest numerical value of the block priority numbers. When a
plurality of block priority numbers each have a numerical value equal to the
most significant value, the block priority number associated with the data
classification block having the lowest numerical address (or, alternatively, the
highest numerical address) will be determined to be the most significant block
25 priority number.

At step 4704, enable signals that indicate whether a data classification block originated the most significant block priority number are generated and transmitted to the first and second data classification blocks.

At step 4705, match signals on the internal address lines that indicate a match and that correspond to the first block priority number are propagated in response to an enable signal that indicates that the most significant block priority number originates from the first data classification block. Match signals on the internal address lines that corresponds to the second block priority number are propagated in response to an enable signal that indicates that the most significant block priority number also originates from the second data classification block. According to an embodiment of the present invention, the internal address lines are propagated by a row enable logic circuit.

At step 4706, a first block index is generated. The first block index references a location within the first data table of the first data classification block where information associated with the first block priority number is stored.

At step 4707, a second block index is generated. The second block index references a location within the second data table of the second data classification block where information associated with the second block priority number is stored.

At step 4708, a first match flag signal is generated indicating whether the most significant block priority number originates from first data classification block and a second match flag signal is generated indicating whether the most significant block priority number originates from the second data classification block. According to an embodiment of the present invention, whether the most significant block priority number originates from a data classification block is

determined from the match signals on the internal address lines from a row enable logic circuit in a data classification block.

At step 4709, a device index is determined from the first and second match flag signals and the first and second block indices.

5 Figure 48 is a flow chart illustrating a method for determining a device index from the first and second match flag signals and the first and second block indices according to an embodiment of the present invention. At step 4801, a most significant block identifier is determined. The most significant block identifier is determined by associating a block identifier of the data classification
10 block generating the most significant block priority number as the most significant block identifier as indicated by the match flag signals. When more than one data classification block originates the most significant block priority number, the data classification block having the lowest numerical address or having highest priority is determined to have the most significant block
15 identifier.

At step 4802, the block indices are stored in an index storage unit. The first block index is stored in a first storage element having a first address. The second block index is stored in a second storage element having a second address.

20 At step 4803, the most significant block identifier is decoded into an address. According to an embodiment of the present invention, the most significant block identifier is decoded into an address of a storage element that stores a block index from a data classification block that corresponds to the most significant block identifier.

25 At step 4804, the most significant block index is output. A block index is output from the index storage unit in response to asserting the decoded address. This block index is designated as being the most significant block index.

At step 4805, a device index is generated. According to an embodiment of the present invention, the device index is generated by designating the most significant block identifier as the most significant bits of the device index and by designating the most significant block index as the least significant bits of the device index.

For an alternative embodiment in which index storage unit 4640 is replaced by a multiplexer or other select circuit, step 4802 may be omitted and steps 4803 and 4804 select and output the appropriate BIN from the multiplexer in response to the decoded MSBID.

Figure 49 illustrates an example of the segmented classification system of Figure 44 generating a device index for the highest priority matching policy statement (or IP address with associated mask) in system 4400. The process of generating a device index is illustrated with the aid of Figure 47. In this example, each data table has at least one entry that matches a search key. At step 4701, data classification block 4310 determines a priority number corresponding to the highest priority matching data stored in the data table 4311 that matches the search key. In this example, the priority logic in the priority index table 4312 determines that the block priority number for data classification block 4310 is 0000. Each entry in data table 4311 having priority number 0000 asserts a match state for match signals on the corresponding internal address lines in the first data classification block 4310.

At step 4702, data classification block 4320 determines a priority number corresponding to the highest priority matching data stored in the data table 4321 that matches the search key. In this example, the priority logic in the priority index table 4322 determines that the block priority number for the data classification block 4320 is also 0000. Each entry in data table 4311 having priority number 0000 asserts a match state for match signals on the

corresponding internal address lines in the second data classification block 4320 associated with the block priority number are asserted.

It should be appreciated that the classification system 600 (shown in Figure 6) may be segmented into any number data classification blocks. In this example, the classification system 600 is segmented into three data classification blocks. Also, at step 4701 or 4702, the third data classification block 4330 determines a priority number corresponding to the highest priority matching data stored in data table 4331 that matches the search key. The priority logic in the priority index table 4332 determines that the block priority number for data classification block 4330 is 0010. Each entry in data table 4311 having priority number 0010 asserts a match state for match signals on the corresponding internal address lines in the third data classification block 4330 associated with the block priority number are asserted.

At step 4703, the most significant block priority number is determined. The most significant value is determined by comparing block priority numbers received from data classification blocks 4310, 4320, and 4330 for the lowest numerical value. The compare logic 4410 compares the block priority numbers and determines that the most significant block priority number is 0000.

At step 4704, enable signals are generated to the data classification blocks indicating whether the data classification block originated the most significant block priority number. The inhibit signal generator 4420 receives the most significant block priority number from the compare logic 4410 and compares the most significant block priority number with the block priority numbers received from each of the data classification blocks 4310, 4320, and 4330. The inhibit signal generator 4420 determines that data classification blocks 4310 and 4320 originate the most significant block priority number, 0000, and transmits enable signals to data classification blocks 4310 and 4320 indicating that they originate

the most significant block priority number. The inhibit signal generator 4420 also determines that data classification block 4330 did not originate the most significant block priority number, 0000, and transmits an enable signal to data classification block 4330 indicating that it did not originate the most significant
5 block priority number. According to an embodiment of the present invention, the enable signals to the data classification blocks 4310 and 4320 may be in the form of an asserted signal line and the enable signal to the data classification block 4330 may be in the form of a non-asserted signal line.

At step 4705, the row enable logic circuit 4313 propagates the match
10 signals on the internal address lines from the priority index table 4312 to the encoder 4315 in response to receiving an enable signal indicating that the most significant block priority number originates from data classification block 4310. Similarly, the row enable logic circuit 4323 propagates the match signals on the internal address lines from the priority index table 4322 to the encoder 4325 in
15 response to receiving an enable signal indicating that the most significant block priority number originates from data classification block 4320. The row enable logic circuit 4333 provides mismatch signals to encoder 4335 in response to receiving an enable signal indicating that the most significant block priority number does not originate from data classification block 4330.

At step 4706, a first block index for the first data classification block 4310
20 is determined. The encoder 4315 determines an index or address corresponding to the most significant priority matching entry in data table 4311 as the first block index for the data classification block 4310. If multiple internal address lines are asserted, the encoder 4315 may select the lowest index associated with
25 the asserted internal address lines as the first block index. The encoder 4315 determines that the block index for data classification block 4310 is 0010. Thus,

the highest priority matching entry for block 4310 is stored at address 2 in data table 4311.

At step 4707, a second block index for the second data classification block 4320 is determined. The encoder 4325 determines an index or address

5 corresponding to the most significant priority matching entry in data table 4321 as the second block index for the data classification block 4320. If multiple internal address lines are asserted, the encoder 4325 may select the lowest index associated with the asserted internal address lines as the second block index.

The encoder 4325 determines that the block index for data classification block
10 4320 is 0100. Thus, the highest priority matching entry for block 4320 is stored at address 4 in data table 4321. Steps 4706 and 4707 may be performed sequentially or simultaneously.

At step 4706 or 4707, a third block index for the third data classification block 4330 is determined. Since the internal address lines from the row enable
15 circuit 4333 indicate mismatch results, the encoder 4335 generates a default index that is designated as that the block index for data classification block 4320. In this example, the default index is the highest (least significant) index associated with the third data classification block 4330, 1111.

At step 4708, match flag signals for the data classification blocks are
20 generated. Match flag generator 4314 asserts MF 0 to a high state indicating that data classification block 4310 originates the most significant block priority number. Match flag generator 4324 asserts MF 1 to a high state indicating that the data classification block 4320 also originates the most significant block priority number. Match flag generator 4334 de-asserts MF N to a low state
25 indicating that the data classification block 4330 does not originate the most significant block priority number.

At step 4709, a device index is determined from the match flag signals and the block indices. According to an embodiment of the present invention, the device index is generated from a most significant block identifier and a most significant block index. A block identifier that identifies the data classification block that originates the most significant block identifier is given the designation of the most significant block identifier. When more than one data classification block originates the most significant block identifier, the block identifier corresponding to the data classification block having the highest priority is designated the most significant block identifier. According to an embodiment of the data classification system 4400, data classification block 4310 has the highest priority, data classification block 4320 has the second highest priority, and data classification block 4330 has the Nth highest priority. In this example, match flag signals indicate that both data classification blocks 4310 and 4320 originate the most significant block priority number. Thus, the block identifier for data classification block 4310, 0000, is given designation of the most significant block identifier. A block index that originates from the data classification block having the most significant block identifier is given the designation of the most significant block index. In this example, the block index 0010 originating from data classification block 4310 is given the designation of the most significant block index. According to an embodiment of the present invention, the address generator 4430 designates the most significant block identifier, 0000, as the most significant bits of the device index and designates the most significant block index, 0010, as the least significant bits of the device index. Thus generating the device index 00000010.

Figure 50 illustrates a segmented data classification system 5000 that includes a device index processor 5060 that is one embodiment of the device index processor 4360 shown in Figure 43. The device index processor 5060

includes a compare logic 5010 and an inhibit signal generator 5020 that is configured to operate in a pipeline mode of operation where the compare logic 5010 may process block priority numbers from more than one group of data at a time. Each priority index table 4312, 4322, and 4332 may be segmented into X
5 equal segments, where X may be any number. Each segment may be separated by a register block (not shown) that registers intermediate compare results on match line segments (not shown) within the priority index tables 4312, 4322, and 4332. The registered results of each segment may be transmitted to the compare logic 5010 as sub-block priority numbers (SBPNs) where the sub-block priority
10 numbers are processed to determine most significant sub-block priority numbers. The compare logic 5010 may receive sub-block priority numbers SBPN0₀-SBPN0_x for a first set of data in the data classification block 4310. The compare logic 5010 may receive sub-block priority numbers SBPN1₀-SBPN1_x for a first set of data in the data classification block 4320. The compare logic 5010
15 may receive sub-block priority numbers SBPNN₀-SBPNN_x for a first set of data in the data classification block 4330.

Figure 51 illustrates a compare logic 5110 that is one embodiment of the compare logic 5010 shown in Figure 50. The compare logic 5010 includes a plurality of compare logic circuits 5120, 5130, and 5140. Each of the compare
20 logic circuits 5120, 5130, and 5140 processes a set of sub-block priority numbers from each of the data classification blocks 4310, 4320, and 4330. It should be appreciated that the number of compare logic circuits in the compare logic 5110 could correspond to the number of segments that the priority index tables 4312, 4322, and 4332 are segmented into and the number of sets of sub-block priority
25 numbers generated by the priority index tables 4312, 4322, and 4332. Compare logic circuit 5120 processes a first set of sub-block priority numbers which represent the most significant bits of the block priority numbers from the data

classification blocks 4310, 4320, and 4330. The compare logic circuit 5120 includes a first stage comparator 5122 that determines a first most significant sub-block priority numbers, MSSBPN₀, from the first set of sub-block priority numbers SBPN₀, SBPN₁, and SBPNN₀.

5 The compare logic circuit 5120 includes a plurality of second stage comparators 5123-5125 that compares the first most significant sub-block priority number with the first set of sub-block priority numbers. Second stage comparator 5123 compares the first most significant sub-block priority number, MSSBPN₀, with the first sub-block priority number, SBPN₀, from data
10 classification block 4310. If the numbers match, second stage comparator 5123 generates an output that operates as a control signal, CNTLA₀, that indicates whether data classification block 4310 originates the first most significant sub-block priority number MSSBPN₀. Second stage comparator 5124 compares the first most significant sub-block priority number, MSSBPN₀, with the first sub-
15 block priority number, SBPN₁, from data classification block 4320. If the numbers match, second stage comparator 5124 generates an output that operates as a control signal, CNTLA₁, that indicates whether data classification block 4320 originates the first most significant sub-block priority number MSSBPN₀. Second stage comparator 5125 compares the first most significant sub-block priority
20 number, MSSBPN₀, with the first sub-block priority number, SBPNN₀, from data classification block 4330. If the numbers match, second stage comparator 5125 generates an output that operates as a control signal, CNTLA_N, that indicates whether data classification block 4330 originates the first most significant sub-block priority number MSSBPN₀.

25 Compare logic circuit 5130 includes a compare enable circuit 5131 and first stage comparator 5132. First stage comparator 5132 determines a second most significant sub-block priority, MSSBPN₁, from the second set of sub-block

priority numbers $SBPN_{0,1}$, $SBPN_{1,1}$, and $SBPNN_1$. The compare enable circuit 5131 receives a second set of sub-block priority numbers $SBPN_{0,1}$, $SBPN_{1,1}$, and $SBPNN_1$ from the data classification blocks 4310, 4320, and 4330 and control signals $CNTLA_1$, $CNTLA_2$, and $CNTLA_N$ from compare logic circuit 5120. If a particular control signal indicates that the corresponding first SBPN (i.e., $SBPN_{0,0}$, $SBPN_{1,0}$, and $SBPNN_0$) is $MSSBPN_0$, compare enable circuit 5131 provides the corresponding second SBPN (i.e., $SBPN_{0,1}$, $SBPN_{1,1}$, and $SBPNN_1$) to first stage comparator 5132 unchanged. If, however, a particular control signal indicates that the corresponding first SBPN is not $MSSBPN_0$, compare enable circuit 5131 assigns a least significant sub-block priority number (LSSBPN) to the corresponding second sub-block priority number from that data classification block. When a sub-block priority number from a data classification block is assigned a least significant sub-block priority number, that data classification block does not originate the most significant sub-block priority number for that corresponding set of sub-block priority numbers.

For one embodiment, compare enable circuit 5131 may be a set of OR logic circuits that each have an input for receiving one bit of an SBPN, and another input for receiving a corresponding CNTLA signal. For other embodiments, any other circuits or logic gates may be used.

The compare logic circuit 5130 includes a plurality of second stage comparators 5133-5135 that compares the second most significant sub-block priority number with the second set of sub-block priority numbers. Second stage comparator 5133 compares the second most significant sub-block priority number, $MSSBPN_1$, with the second sub-block priority number, $SBPN_{0,1}$, from data classification block 4310. If the numbers match, second stage comparator 5133 generates a signal that indicates that data classification block 4310 originates the second most significant sub-block priority number $MSSBPN_1$.

Second stage comparator 5134 compares the second most significant sub-block priority number, MSSBPN₁, with the second sub-block priority number, SBPN₁, from data classification block 4320. If the numbers match, second stage comparator 5134 generates a signal that indicates that data classification block 4320 originates the second most significant sub-block priority number MSSBPN₁. Second stage comparator 5135 compares the second most significant sub-block priority number, MSSBPN₁, with the second sub-block priority number, SBPN₁, from data classification block 4330. If the numbers match, second stage comparator 5135 generates a signal that indicates that data classification block 4330 originates the second most significant sub-block priority number MSSBPN₁.

The compare logic circuit 5130 includes a plurality of control signal processors 5136, 5137, and 5138. Control signal processor 5136 receives control signal CNTLA₀ from second stage comparator 5123 and the signal output from second stage comparator 5133. Control signal processor 5136 generates a control signal, CNTLB₀, that indicates whether data classification block 4310 originates both the first and second most significant sub-block priority numbers, MSSBPN₀ and MSSBPN₁. Control signal processor 5137 receives control signal CNTLA₁ from second stage comparator 5124 and the signal output from second stage comparator 5134. Control signal processor 5137 generates a control signal, CNTLB₁, that indicates whether data classification block 4320 originates both the first and second most significant sub-block priority numbers, MSSBPN₀ and MSSBPN₁. Control signal processor 5138 receives control signal CNTLA_N from second stage comparator 5125 and the signal output from second stage comparator 5135. Control signal processor 5138 generates a control signal, CNTLB_N, that indicates whether data classification block 4330 originates both the first and second most significant sub-block priority numbers, MSSBPN₀ and MSSBPN₁. For one embodiment, the control signal processors may be AND logic

circuits that logically combine CNTLA and signals output by the second stage comparators 5133-5135 to generate control signals CNTLB. For other embodiments, the control signal processors may be one or more logic gates or other circuitry.

5 Compare logic circuit 5140 includes a compare enable circuit 5141 and first stage comparator 5142. First stage comparator 5142 determines a third most significant sub-block priority, $MSSBPN_x$, from the third set of sub-block priority numbers $SBPN0_x$, $SBPN1_x$, and $SBPNN_x$. The compare enable circuit 5141 receives a third set of sub-block priority numbers $SBPN0_x$, $SBPN1_x$, and $SBPNN_x$
10 from the data classification blocks 4310, 4320, and 4330 and control signals $CNTLB_0$, $CNTLB_1$, and $CNTLB_N$ from compare logic circuit 5130. If a particular control signal indicates that the corresponding second SBPN (i.e., $SBPN0_1$, $SBPN1_1$, and $SBPNN_1$) is $MSSBPN_1$, compare enable circuit 5141 provides the corresponding third SBPN (i.e., $SBPN0_2$, $SBPN1_2$, and $SBPNN_2$) to first stage
15 comparator 5142 unchanged. If, however, a particular control signal indicates that the corresponding first SBPN is not $MSSBPN_1$, compare enable circuit 5141 assigns a least significant sub-block priority number to the corresponding third sub-block priority number from that data classification block. When a sub-block priority number from a data classification block is assigned a least significant
20 sub-block priority number, that data classification block does not originate the most significant sub-block priority number for that corresponding set of sub-block priority numbers.

 The compare logic circuit 5140 includes a plurality of second stage comparators 5143-5145 that compares the third most significant sub-block
25 priority number with the third set of sub-block priority numbers. Second stage comparator 5143 compares the third most significant sub-block priority number, $MSSBPN_x$, with the third sub-block priority number, $SBPN0_x$, from data

classification block 4310. If the numbers match, second stage comparator 5143 generates a signal that indicates that data classification block 4310 originates the third most significant sub-block priority number $MSSBPN_X$. Second stage comparator 5144 compares the third most significant sub-block priority number, $MSSBPN_X$, with the third sub-block priority number, $SBPN1_X$, from data classification block 4320. If the numbers match, second stage comparator 5144 generates a signal that indicates that data classification block 4320 originates the third most significant sub-block priority number $MSSBPN_X$. Second stage comparator 5145 compares the third most significant sub-block priority number, $MSSBPN_X$, with the third sub-block priority number, $SBPNN_X$, from data classification block 4330. If the numbers match, second stage comparator 5145 generates a signal that indicates that data classification block 4330 originates the third most significant sub-block priority number $MSSBPN_X$.

The compare logic circuit 5140 includes a plurality of control signal processors 5146, 5147, and 5148. Control signal processor 5146 receives control signal $CNTLB_0$ from control signal processor 5136 and the signal output from second stage comparator 5143. Control signal processor 5146 generates a control signal, $CNTLX_0$, that indicates whether data classification block 4310 originates the first, second, and third most significant sub-block priority numbers, $MSSBPN_0$, $MSSBPN_1$, and $MSSBPN_X$. Control signal processor 5147 receives control signal $CNTLB_1$ from control signal processor 5137 and the signal output from second stage comparator 5144. Control signal processor 5147 generates a control signal, $CNTLX_1$, that indicates whether data classification block 4320 originates the first, second, and third most significant sub-block priority numbers, $MSSBPN_0$, $MSSBPN_1$, and $MSSBPN_X$. Control signal processor 5148 receives control signal $CNTLB_N$ from control signal processor 5138 and the signal output from second stage comparator 5145. Control signal processor 5148

generates a control signal, $CNTLX_N$, that indicates whether data classification block 4330 originates the first, second, and third most significant sub-block priority numbers, $MSSBPN_0$, $MSSBPN_1$, and $MSSBPN_X$.

According to an embodiment of the compare logic 5110, each of the
5 second stage comparators 5123-5125, 5133-5135, and 5143-5145 may be implemented using a plurality of circuits that perform an XNOR function. The output of the circuits that perform the XNOR function may be input into a circuit that performs an AND function. According to an embodiment of the compare logic 5110, each of the control signal processors 5136-5138, and 5146-5148 may be
10 implemented using a circuit that perform an AND function. It should be appreciated that the second stage comparators 5123-5125, 5133-5135, and 5143-5145, and the control signal processors 5136-5138, and 5146-5148 may be implemented using any known circuitry or technique.

It should be appreciated that each of the compare logic circuits 5120, 5130,
15 and 5140 may be implemented using a priority index table as shown in Figure 52. Figure 52 illustrates an alternate embodiment of the compare logic 5010 shown in Figure 50. Each of the compare logic circuits 5220, 5230, and 5240 are implemented using a priority index table that includes a plurality of memory element/priority logic element pair rows (not shown) that store and compare the
20 bits of sub-block priority numbers. Each memory element/priority logic element pair row in a compare circuit is coupled to a corresponding memory element/priority logic element pair row in a same relative position in an adjacent compare circuit via a match line segment. The sub-block priority numbers from a data classification block are stored in memory element/priority
25 logic element pair rows of the same relative position in the compare circuits 5220, 5230, and 5240.

Figure 53 illustrates an inhibit signal generator 5320 that is an embodiment of the inhibit signal generator 5020 shown in Figure 50. The inhibit signal generator 5320 includes synchronizing logic 5310 that operates to synchronize the control signals received from the compare circuits 5120, 5130, and 5140 shown in Figure 51 (or 5220, 5230, and 5240 shown in Figure 52). In one embodiment, the synchronizing logic 5310 includes two register blocks for storing the control signals received from compare circuit 5120 (or 5220) for two clock cycles, and a single register block for storing the control signals received from compare circuit 5130 (or 5230) for one clock cycle. This allows the control signals received from compare circuits 5120 and 5130 (or 5220 and 5230) to be synchronized with control signals received from compare circuit 5140 (or 5240). It will be appreciated that any other number of register blocks may be used depending on how many stages each priority index table is segmented into.

The inhibit signal generator 5320 also includes an enable signal generator 5330. The enable signal generator 5330 generates a first enable signal, EN 0, to the data classification block 4310 if the control signals CNTLA₀, CNTLB₀, and CNTLX₀ indicate that the data classification block 4310 originates the first, second, and third most significant sub-block priority numbers. The enable signal generator 5330 generates a second enable signal, EN 1, to the data classification block 4320 if the control signals CNTLA₁, CNTLB₁, and CNTLX₁ indicate that the data classification block 4320 originates the first, second, and third most significant sub-block priority numbers. The enable signal generator 5330 generates a third enable signal, EN N, to the data classification block 4330 if the control signals CNTLA_N, CNTLB_N, and CNTLX_N indicate that the data classification block 4330 originates the first, second, and third most significant sub-block priority numbers. According to an embodiment of the inhibit signal generator 5320, the enable signal generator 5330 is implemented using a

plurality of OR logic circuits. For example, each OR logic circuit may combine all of the CNTLA₀, CNTLB₀, etc. to generate EN₀; another OR logic circuit may combine all of CNTLA₁, CNTLB₁, etc. to generate EN₁, and so on while another OR logic circuit may combine all of CNTLA_N, CNTLB_N, etc. to generate EN_N. It should be appreciated that the inhibit signal generator 5330 may be implemented using any known circuitry or technique.

In a first alternate embodiment of the inhibit signal generator 5320, the synchronization logic 5310 may be implemented in the compare logic 5010 (shown in Figure 50). In a second alternate embodiment of the inhibit signal generator 5320, the synchronization logic 5310 is omitted. In this embodiment, the enable signal generator 5330 is implemented with a first connector connecting control signal CNTLX₀ to the data classification block 4310 as an enable signal EN 0, a second connector connecting control signal CNTLX₁ to the data classification block 4320 as enable signal EN 1, and a third connector connecting control signal CNTLX_N to the data classification block 4330 as an enable signal EN N.

According to an embodiment of the segmented data classification system 5000 (shown in Figure 50), the data classification blocks 4310, 4320, and 4330, and the device index processor 5060 all reside on a single semiconductor substrate.

Figure 54 is a flow chart illustrating a method for generating a device index in a segmented data classification system according to the present invention. The steps illustrated in Figure 54 may be used in place of steps 4701-4704 of Figure 47 when the segmented data classification system is operating in a pipelined mode. At step 5401, a search key is compared to valid entries in data tables of data classification blocks to identify matching entries.

At step 5402, first segments in the priority index tables of each data classification block compare a first plurality of bits of stored priority numbers associated with corresponding matching entries in its data table. The first plurality of bits of the stored priority number having the most significant value
5 in each priority index table is designated as the first sub-block priority number for its corresponding data classification block.

At step 5403, a first most significant sub-block priority number is determined from a first sub-block priority number from a first data classification block and a first sub-block priority number from a second data classification
10 block.

At step 5404, the data classification block that originates the first most significant sub-block priority number is determined. According to an embodiment of the present invention, the determination is made by comparing the first sub-block priority numbers from the first and second data classification
15 blocks with the first most significant sub-block priority number to find a match.

At step 5405, a first control signal that indicates whether the first most significant sub-block priority number originates from the first data classification block is generated.

At step 5406, a second control signal that indicates whether the first most
20 significant sub-block priority number originates from the second data classification block is generated.

At step 5407, second segments in the priority index tables of each data classification block compare a second plurality of bits of stored priority numbers
25 associated with the corresponding matching entries in its data table. The second plurality of bits of the stored priority number having the most significant value

in each priority index table is designated as the second sub-block priority number for its corresponding data classification block.

At step 5408, a second most significant sub-block priority number is determined from the second sub-block priority numbers from the first and
5 second data classification blocks.

At step 5409, the data classification block that originates the second most significant sub-block priority number is determined. According to an embodiment of the present invention, the determination is made by comparing the second sub-block priority numbers from the first and second data
10 classification blocks with the second most significant sub-block priority number to find a match.

At step 5410, a third control signal that indicates whether the first and second most significant sub-block priority number originates from the first data classification block is generated.

At step 5411, a fourth control signal that indicates whether the first and second most significant sub-block priority number originates from the second data classification block is generated.
15

At step 5412, the control signals are synchronized. According to an embodiment of the present invention, control signals that are generated earlier
20 are registered until all the control signals have been generated.

At step 5413, enable signals that indicate whether a data classification block originates all the most significant sub-block priority numbers are generated. According to an embodiment of the present invention, a first enable signal for the first data classification block is generated in response to the first
25 and third control signals, and a second enable signal for the second data classification block is generated in response to the second and fourth control signals. According to an alternate embodiment of the present invention, steps

5412 is omitted and the first enable signal is generated in response to the third control signal and the second enable signal is generated in response to the fourth control signal. It should be appreciated that steps similar to the steps described in steps 5407-5411 may be repeated if additional sub-block priority numbers
5 need to be processed.

Figures 47, 48, and 54 illustrate flow charts describing methods for priority encoding and generating a device index according to embodiments of the present invention. Some of the steps illustrated in these figures may be performed in an order other than that which is described. It should be
10 appreciated that not all of the steps described are required to be performed and that some of the illustrated steps may be substituted with other steps.

The pipeline compare operation of the system described in Figures 50 and 51 is further described with the aid of the qualitative timing diagram of Figure 55. In this example, each priority index table 4312, 4322, and 4332 is segmented
15 into three equal segments where each segment stores and compares an 8 bit segment of a 24 bit priority number. In the first clock cycle, a first compare operation COMPARE1 is provided to each data table 4311, 4321, and 4331 and a first search key compared with all of the valid entries. The search or compare results are provided to the priority index tables 4312, 4322, and 4332. During the
20 second (or later) clock cycle, the first segments of each priority index table compare 8 bits of the stored priority numbers associated with corresponding matching entries to determine the most significant bits of the block priority numbers associated with each index table; namely, SBPN₀, SBPN₁, and SBPNN₀, respectively. SBPN₀, SBPN₁, and SBPNN₀ are provided to compare
25 logic circuit 5120. The compare logic circuit 5120 determines the first most significant sub-block priority number, MSSBPN₀, determines which data classification block originates the first most significant sub-block priority

number, $MSSBPN_0$, and generates control signals that indicate whether a data classification block originates the first most significant sub-block priority number, $MSSBPN_0$ during the third clock cycle. Also during the third (or later) clock cycle, the second segments of the priority index tables 4312, 4322, and 4332

5 compare the next 8 bits of their stored priority numbers associated with corresponding matching entries to determine $SBPN0_1$, $SBPN1_1$, and $SBPNN_1$, respectively. $SBPN0_1$, $SBPN1_1$, and $SBPNN_1$ are provided to the compare logic circuit 5130. The compare logic circuit 5130 then determines the second most significant sub-block priority number $MSSBPN_1$, determines which data

10 classification block originates the second most significant block priority number, $MSSBPN_1$, and generates control signals that indicate whether a data classification block originates the first and second most significant block priority number, $MSSBPN_0$ and $MSSBPN_1$, during the fourth clock cycle. Similarly, during the fourth (or later) clock cycle, the third segments of the priority index

15 tables 4312, 4322, and 4332 compare the last 8 bits of their stored priority numbers associated with corresponding matching entries to determine $SBPN0_x$, $SBPN1_x$, and $SBPNN_x$, respectively. $SBPN0_x$, $SBPN1_x$, and $SBPNN_x$ are provided to compare logic circuit 5140. The compare logic circuit 5140 then determines the third most significant sub-block priority number, $MSSBPN_x$,

20 determines which data classification block originates the third most significant sub-block priority numbers, $MSSBPN_x$, and generates control signals that indicate whether a data classification block originates the first, second and third most significant block priority numbers, $MSSBPN_0$, $MSSBPN_1$, and $MSSBPN_x$, during the fifth clock cycle. Also during the fifth clock cycle, the inhibit signal

25 generator 5020 generates enable signals to the row enable logic 4313, 4323, and 4333 in response to the control signals from the compare logic circuit 5140. During the sixth clock cycle, BINs for each of the data classification blocks 4310,

4320, and 4330 are generated by encoders 4315, 4325, and 4335 in response to signals on internal address lines from the row enable logic 4313, 4323, and 4333. Also during the sixth clock cycle, a DIN that corresponds to the first search key is generated in response to the BINs and match flag signals from each of the data
5 classification blocks 4310, 4320, and 4330.

A second compare operation can be pipelined with the first compare operation also as shown in Figure 55. During the second clock cycle, a second compare operation, COMPARE2 may be provided to each data table 4311, 4321, and 4331 and a second search key compared with all valid entries. The search or
10 compare results for COMPARE2 are provided to the first segments in the priority index tables 4312, 4322, and 4332. During the third (or later) clock cycle, the first segments of each priority index table compare 8 bits of the stored priority numbers associated with corresponding matching entries to determine the most significant bits of the block priority numbers associated with each index
15 table, SBPN0₀, SBPN1₀, and SBPNN₀, respectively. SBPN0₀, SBPN1₀, and SBPNN₀ are provided to compare logic circuit 5120. The compare logic circuit 5120 determines the first most significant block priority number, MSSBPN₀, determines which data classification block originates the first most significant block priority number, MSSBPN₀, and generates control signals that indicate
20 whether a data classification block originates the first most significant block priority number, MSSBPN₀ during the fourth clock cycle. Also during the fourth (or later) clock cycle, the second segments of the priority index tables 4312, 4322, and 4332 compare the next 8 bits of their stored priority numbers associated with corresponding matching entries to determine SBPN0₁, SBPN1₁, and SBPNN₁,
25 respectively. SBPN0₁, SBPN1₁, and SBPNN₁ are provided to the compare logic circuit 5130. The compare logic circuit 5130 then determines the second most significant sub-block priority number MSSBPN₁, determines which data

classification block originates the second most significant block priority number, MSSBPN₁, and generates control signals that indicate whether a data classification block originates the first and second most significant block priority number, MSSBPN₀ and MSSBPN₁, during the fifth clock cycle. Similarly, during the fifth (or later) clock cycle, the third segments of the priority index tables 4312, 4322, and 4332 compare the last 8 bits of their stored priority numbers associated with corresponding matching entries to determine SBPN0_x, SBPN1_x, and SBPNN_x, respectively. SBPN0_x, SBPN1_x, and SBPNN_x are provided to compare logic circuit 5140. The compare logic circuit 5140 then determines the third most significant sub-block priority number, MSSBPN_x, determines which data classification block originates the third most significant sub-block priority numbers, MSSBPN_x, and generates control signals that indicate whether a data classification block originates the first, second and third most significant block priority numbers, MSSBPN₀, MSSBPN₁, and MSSBPN_x, during the sixth clock cycle. Also during the sixth clock cycle, the inhibit signal generator 5020 generates enable signals to the row enable logic 4313, 4323, and 4333 in response to the control signals from the compare logic circuit 5140. During the seventh clock cycle, BINs for each of the data classification blocks 4310, 4320, and 4330 are generated by encoders 4315, 4325, and 4335 in response to signals on internal address lines from the row enable logic 4313, 4323, and 4333. Also during the seventh clock cycle, a DIN that corresponds to the second search key is generated in response to the BINs and match flag signals from each of the data classification blocks 4310, 4320, and 4330.

The segmented data classification systems shown in Figures 43, 44, 49, and 50 have been described for classifying policy statements. It should be appreciated that the segmented data classification systems described can also be used to process IP packets that use the CIDR scheme. The segmented data

classification system described can be used to process CIDR based IP addresses without adding additional logic or signal lines to the segmented data classification system. IP addresses and corresponding decoded prefix data can be loaded into the data tables 4311, 4321, and 4331 of each of the data

5 classification blocks 4310, 4320, and 4330. Prefixes of the IP addresses may be stored in the priority index tables 4312, 4322, and 4332 of the data classification blocks 4310, 4320, and 4330. Each of the data classification blocks 4310, 4320, and 4330 would output a block priority number that corresponds to a prefix of a best matched IP address in a data classification block. The device index
 10 processors described in Figures 43, 44, 49, and 50 select the longest prefix mask (i.e., the most significant priority number) from block prefix masks (i.e., the block priority numbers) and generate a device index that may be used to access route information corresponding to the best matched (i.e., longest prefix match) IP address in the segmented data classification system. Any of the embodiments
 15 of the segmented data classification systems described above can be used to process CIDR based IP addresses and their prefix numbers.

In the foregoing specification the invention has been described with reference to specific exemplary embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without
 20 departing from the broader spirit and scope of the invention. The specification and drawings are, accordingly, to be regarded in an illustrative rather than restrictive sense.